CENTRE FOR
SYSTEMS
SCIENCE

LCCR

Simon Fraser University, Burnaby, B.C., Canada V5A 1S6 / (604) 291-4704

Met*: A Method for Discriminating Metonymy and Metaphor by Computer

Dan Fass

CSS/LCCR TR 89-15

November 1989

92-05698/

Met*: A Method for Discriminating Metonymy and Metaphor by Computer

Dan Fass

Centre for Systems Science
Simon Fraser University
Burnaby, British Columbia, Canada V5A 1S6.
CSNET: fass@cs.sfu.ca

Abstract

The met* method distinguishes selected examples of metonymy from metaphor and from literalness and anomaly in short English sentences. In the met* method, literalness is distinguished from the others because they are all non-literal. Metonymy is discriminated from metaphor and anomaly in a way that (1) supports Lakoff and Johnson's (1980) view that in metonymy one entity stands for another whereas in metaphor one entity is viewed as another, (2) permits chains of metonymies (Reddy, 1979), and (3) allows metonymies to co-occur with instances of either literalness, metaphor, or anomaly. Metaphor is distinguished from anomaly because the former contains a relevant analogy, unlike the latter. The met* method is part of Collative Semantics, a semantics for natural language processing, and has been implemented in a computer program called meta5. Some examples of meta5's analysis of metaphor and metonymy are given. The met* method is compared with approaches from philosophy, linguistics, psychology, and artificial intelligence.

1. Introduction

Most of the considerable literature on metaphor and the smaller one on metonymy (see Van Noppen, De Knop, & Jongen, 1985; Shibles, 1971) consists of philosophical, linguistic and psychological views of them. In artificial intelligence (AI), there are not many computational approaches to either metonymy or metaphor and most of those are paper implementations which have not been coded up and run on a computer, whereas the met* method is an AI method that has.

The met* (pronounced "met-star") method, which is described in section 2, takes its name from a remark made by Yorick Wilks. He used met* to refer collectively to metonymy and metaphor: "*" is a match-anything symbol in the Unix operating system; hence, the token "met*" would match the two tokens "metonymy" and "metaphor." The method is part of Collative Semantics (hereafter CS), which is a semantics for natural language processing. CS, and hence the met* method, has been implemented in a program called meta5 (so-called because it does more than metaphor). Some deficiencies in the method and some possible remedies are discussed. In section 3, the method is related to some of the main philosophical, linguistic and psychological views of metonymy and metaphor, and in section 4, to some other computational approaches.

2. The Met* Method

2.1. Linguistic Background

I have argued elsewhere (Fass, 1989a) that understanding natural language (or semantic interpretation) can be viewed as the integration of *constraints* from inside and outside language. Some within-language constraints are syntactic, others are semantic. Some within-language constraints are *lexical* constraints, that is, constraints possessed by words. Lexical syntactic constraints include those on word order, number and tense. This section describes three *lexical semantic constraints: preferences, assertions* and a lexical notion of *relevance*.

Preferences (Wilks, 1973) are also known as selection restrictions (Katz, 1964) and expectations (Schank, 1975). Preferences are restrictions possessed by word senses of certain parts of speech about the semantic classes of words with which they co-occur, hence an adjective sense has a preference for the semantic class of nouns with which it co-occurs, and a verb sense has preferences for the semantic classes of nouns that fill its case roles. For example, the main sense of the verb 'drink' prefers an animal to fill its agent case role, i.e., it is animals that drink.

The assertion of semantic information was noted by Lees (1960) in the formation of noun phrases and later developed by Katz (1964) as the process of "attribution." Assertions contain information which is possessed by word senses of certain parts of speech and which is imposed onto word senses of other parts of speech, e.g., the adjective 'female' contains information that any noun over which it predicates is of the female sex.

Lexical syntactic and semantic constraints are enforced at certain places in sentences which I call *dependencies*. Within a dependency, the word whose constraints are enforced is called the *source* and the other word the *target* (after Martin, 1985). Syntactic dependencies consist of pairs of words of certain parts of speech in which the source, a word from one part of speech, applies one or more syntactic constraints to the target, another word. Examples of source-target pairs include a determiner and a noun, an adjective and a noun, a noun and a verb, and an adverb and a verb.

(1) "The ship ploughed the waves."

Semantic dependencies occur in the same places as syntactic dependencies. The (metaphorical) sentence (1) contains four semantic dependencies: between the determiner 'the' and the noun 'ship', between 'ship' and the verb stem 'plough', between 'the' and the noun 'waves', and between 'waves' and 'plough'. In each semantic dependency, one word acts as the source and predicates upon the other word, which acts as the target. In (1), 'the' and 'plough' both predicate upon 'ship', and 'the' and 'plough' predicate on 'waves'. Semantic dependencies exist between not just pairs of words but also pairs of word *senses*. For example, the metaphorical reading of (1) is because 'waves' is understood as being the sense meaning "movement of water," not for example the sense meaning "movement of the hand."

Semantic relations result from evaluating lexical semantic constraints in sentences. Every semantic relation has a source (a word whose semantic constraints are applied) and a target (a word which received those constraints). Other terms used to refer to the source and target in a semantic relation include: vehicle and tenor (Richards, 1936), subsidiary subject and principal subject (Black, 1962), figurative term and literal term (Perrine, 1971), referent and subject (Tversky, 1977), secondary subject and primary subject (Black, 1979), source and destination (Winston, 1980), old domain and new domain (Hobbs, 1983a), and base and target (Gentner, 1983).

In CS, seven kinds of semantic relation are distinguished: literal, metonymic, metaphorical, anomalous, redundant, inconsistent, and novel relations (this list may not be exhaustive – there could be others). Combinations of these seven semantic relations are the basis of (at minimum) literalness, metonymy, metaphor, anomaly, redundancy, contradiction, contrariness, and novelty. Semantic relations belong to two classes, the *preference-based* and *assertion-based* classes of relation, depending on the kind of lexical semantic constraint enforced. The preference-based class of semantic relations contains literal, metonymic, metaphorical, and anomalous semantic relations. Literal relations are signified by a satisfied preference whereas the remaining three relations, which are all non-literal, are indicated by a violated preference.

(2) "The man drank beer."

There is a literal relation between 'man' and 'drink' in (2) because 'drink' prefers an animal as its agent and a man is a type of animal so the preference is satisfied.

In a metonymy, one entity is used to refer to another that is related to it (Lakoff & Johnson, 1980, p. 35). Several attempts have been made to organize instances of metonymy into categories (e.g., Lakoff & Johnson, 1980; Stern, 1931; Yamanashi, 1987) or "metonymic concepts," as Lakoff and Johnson call them. A common example is PART FOR WHOLE, otherwise known as synechdoche.

(3) "Dave drank the *glasses*" (= the liquid in the glasses).

(4) "The *kettle* is boiling" (= the liquid in the kettle)
(Waldron, 1967, p. 186; Yamanashi, 1987, p. 78).

CONTAINER FOR CONTENTS, another metonymic concept, occurs in (3) between 'drink' and the sense of 'glasses' meaning "containers," and also in (4). In (3), 'drink' has an object preference for a potable liquid, but there is a preference violation because glasses are not potable liquids. It is not glasses that are drunk, but the potable *liquids* in them. There is a relationship here between a CONTAINER (a glass) and its typical CONTENTS (a liquid): this relationship is the metonymic concept CONTAINER FOR CONTENTS.

(5) "You'll find better ideas than that in the library" (Reddy, 1979, p. 309).

Reddy (1979) has observed that metonymies can occur in chains. He suggests that (5) contains a chain of PART FOR WHOLE metonymies between 'ideas' and 'library': the ideas are expressed in words, words are printed on pages, pages are in books, and books are found in a library.

Metonymy and metaphor, in my view, differ because the core of a metonymy is a chain of semantic relationships like CONTAINER FOR CONTENTS whereas the core of a metaphor is a relevant analogy. No one else (to my knowledge) has emphasized the role of relevance in the discovery of an analogy central to a metaphor, though the importance of relevance in recognizing metaphors has been discussed by, e.g., Hobbs (1983b) and Tversky (1977), and the centrality of some analogy or correspondence in metaphorical relations has been made by Gentner (1983, p. 162), Ortony (1979, p. 167), Tourangeau and Sternberg (1982, pp. 218-221) and Wilks (1978), among others.

(6) "The car drank gasoline" (adapted from Wilks, 1978).

The form of relevance used is a lexical notion – i.e., the third kind of lexical semantic constraint – that what is relevant in a sentence is given by the sense of the main sentence verb being currently analyzed. Thus, it is claimed that the semantic relation between 'car' and 'drink' in (6) is metaphorical because there is a preference violation and an underlying relevant analogy between 'car' and 'animal', the preferred agent of 'drink'. A car is not a type of animal, hence the preference violation. However, what is relevant in (6) is drinking, and there is a relevant analogy that animals and cars both use up a liquid of some kind: animals drink potable liquids while cars use gasoline. Hence the metaphorical relation between 'car' and 'drink'. This example will be described in more detail in section 2.3. Incidentally, I do not regard

metaphor as being exclusively recognized by a preference violation but rather that metaphor consists of a contextual constraint violation such as preference violations at the lexical level. This is discussed at more length in section 2.4.

(7) "The idea drank the heart."

Anomalous relations have neither the semantic relationships of a metonymic relation nor the relevant analogy of a metaphorical relation. Hence the semantic relation between 'idea' and 'drink' is anomalous in (7) because 'idea' is not a preferred agent of 'drink' and no metonymic link or relevant analogy can be found between animals (the preferred agent) and ideas.

2.2. Collative Semantics

CS is a semantics for natural language processing that extends many of the main ideas behind Preference Semantics (Wilks, 1973, 1975a, 1975b, 1978). CS has four components: *sense-frames*, *collation*, *semantic vectors*, and *screening*. The met* method is part of the process of collation. Fuller and more general descriptions of the four components appears in Fass (1988a, 1989b).

Sense-frames are dictionary entries for individual word senses. Sense-frames are composed of other word senses that have their own sense-frames, much like Quillian's (1967) planes. Each sense-frame consists of two parts, an *arcs* section and a *node* section, that correspond to the genus and differentia commonly found in a real dictionary definition.

The arcs part of a sense-frame contains a labelled arc to its genus term (a word sense with its own sense-frame). Together, the arcs of all the sense-frames comprise a densely structured semantic network of word senses called the *sense-network*. The node part of a sense-frame is the differentia that provides a "definition" of the word sense represented by the sense-frame that differentiates it from other word senses. The two lexical semantic constraints mentioned earlier, preferences and assertions, play a prominent part in sense-frame nodes.

Sense-frame nodes for nouns (node-type 0) resemble Wilks' (1978) pseudo-texts. The nodes contain lists of two-element and three-element lists called *cells*. Cells contain word senses and have a syntax modelled on English. Each cell expresses a piece of functional or structural information and can be thought of as a complex semantic feature or property of a noun. Figure 1 shows sense-frames for two senses of the noun 'crook'. Crook1 is the sense meaning "thief" and crook2 is the shepherd's tool.

```
sf(crook1,                              sf(crook2,
   [[arcs,                                 [[arcs,
    [[supertype,   criminal1]]],            [[supertype,   stick1]]],
    [node0,                                 [node0,
    [[it1,    steal1,   valuables1]]]]).    [[shepherd1,   use1,   it1],
                                             [it1,    shepherd1,   sheep1]]]]).
```

Figure 1. Sense-frames for crook1 and crook2.

All the terms in sense-frames are word senses with their own sense-frames or words used in a particular sense that could be replaced by word senses. It1 refers to the word sense being defined by the sense-frame so, for example, crook1 can be substituted for it1 in [it1, steal1, valuables1]. Common dictionary practice is followed in that word senses are listed separately for each part of speech and numbered by frequency of occurrence. Hence in crook2, the cell [shepherd1, use1, it1] contains the noun sense shepherd1 while the cell [it1, shepherd1, sheep1] contains the verb sense shepherd1.

Sense-frame nodes for adjectives, adverbs and other modifiers (node-type 1) contain preferences and assertions but space does not permit a description of them here. Sense-frame nodes for verbs and prepositions (node-type 2) are case frames containing case subparts filled by case roles such as 'agent', 'object' and 'instrument'. Case subparts contain preferences, and assertions if the verb describes a state change. Figure 2 shows the sense-frames for the verb senses eat1 and drink1. In both, the agent preference is for an animal but the object preferences differ: the preference of eat1 is for food1, i.e., an edible solid, while the preference of drink1 is for drink1 (the noun sense), i.e., a potable liquid.

```
sf(eat1,                                    sf(drink1,
   [[arcs,                                     [[arcs,
    [[supertype, [ingest1, expend1]]]],         [[supertype, [ingest1, expend1]]]]
    [node2,                                     [node2,
     [[agent,                                    [[agent,
        [preference, animal1]],                     [preference, animal1]],
      [object,                                    [object,
       [preference, food1]]]]]).                   [preference, drink1]]]]]).
```

Figure 2. Sense-frames for eat1 and drink1 (verb).

The second component of CS is the process of collation. Collation matches the sense-frames of two word senses and finds a system of multiple mappings between those sense-frames, thereby discriminating the semantic relations between the word senses. The basic mappings are paths found by a graph search algorithm that operates over the sense-network. Five types of network path are distinguished. Two kinds of network path denote kinds of "inclusion;" three kinds denote "exclusion." These paths are used to build more complex mappings found by a frame-matching algorithm.

The frame-matching algorithm matches the sets of cells from two sense-frames. Seven types of cell match are distinguished. The sets of cells, which need not be ordered, are inherited down the sense-network. A series of structural constraints isolate pairs of cells that are matched using the graph search algorithm. Paths denoting inclusion are sought between terms occupying identical positions in those cells.

A kind of lexical relevance – the third kind of lexical semantic constraint mentioned previously – is found dynamically from the sentence context. Relevance divides the set of cells from the source sense-frame into two subsets. One cell is selected as *relevant* given the context; the remaining cells are termed *non-relevant*. Collation matches both the relevant and non-relevant cells against the cells from the target sense-frame. The notion of relevance is used to distinguish metaphorical from anomalous relations; it is also used when finding CO-AGENT FOR ACTIVITY metonymies.

Figure 3 shows the met* method laid out as a flow chart and illustrates how the preference-based class of semantic relations is discriminated. A preference is evaluated and if it is satisfied (i.e., an inclusive sense-network path is found) then the semantic relation is recognized as a literal one. If the preference is not satisfied (i.e., if the path is not inclusive) a metonymy is sought.

Knowledge of metonymic concepts is encoded in *metonymic inference rules*. The process that finds metonymies is called *metonymic inferencing* – the application of these inference rules. If no successful rule is found then no metonymy is discovered and a metaphorical or anomalous semantic relation is then sought. Conversely, a successful rule means a metonymy has been recognized. A successful rule establishes a relationship between the original source or the target ("one entity") and a term ("another that is related to it") that "stands for" one of them. Depending on the metonymic concept, a successful inference is sometimes directed "forward" from the preference or "backward" from the target (more on this shortly). It is this direction of inferencing that determines whether the source or target is substituted in a successful metonymy. The substitute source or target is used to discover another semantic relation which can be literal, metonymic again, metaphorical, or anomalous. Hence a preference-based semantic relation can consist of (1) a single literal, metaphorical, or anomalous relation or (2) a *multi-relation* comprised of a literal, metaphorical, or anomalous relation plus either a single metonymy or a chain of metonymies. All these combinations, but only these, are derivable from Figure 3.

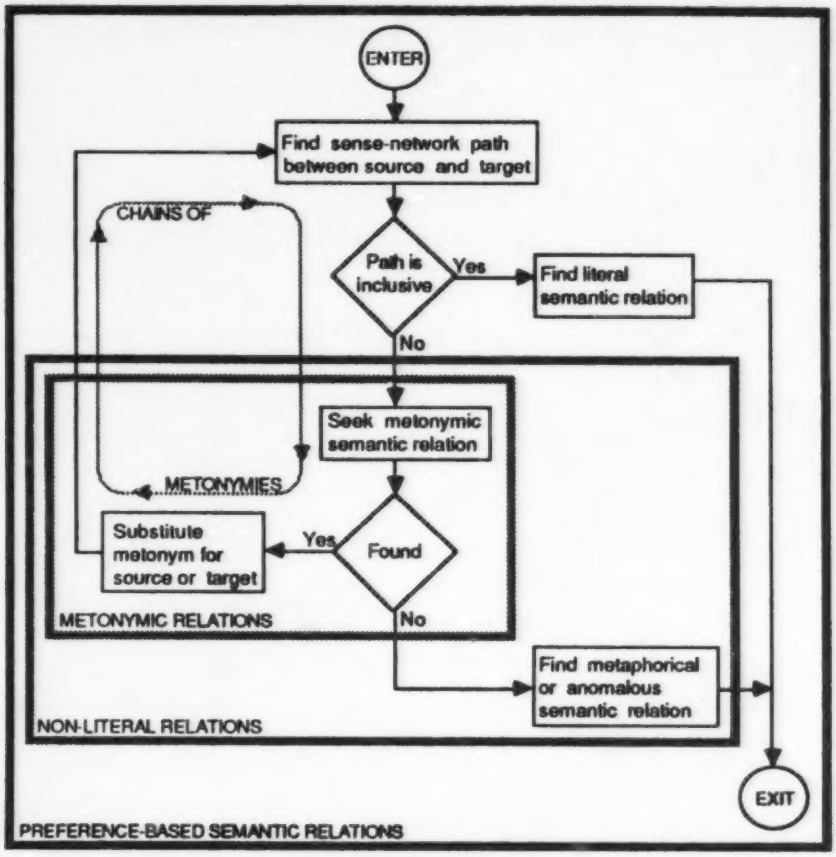

Figure 3. **The met* method.**

Note that in the met* method as presented in Figure 3, semantic relations are tried in a certain order: literal, metonymic, metaphorical, and finally anomalous. This ordering implies that a literal interpretation is sought *before* a non-literal one (cf. Harris, 1976). The ordering results from thinking about discriminating the semantic relations in serial processing terms rather than parallel processing terms, particularly the serial order in which sense-network paths and metonymic inference rules are tried: inclusive network paths (indicating literalness) then metonymic inference rules (metonymy) then non-inclusive network paths (metaphor or anomaly). The serial order and the topology of the flow chart is mainly due to the need to satisfy a number of observations about the preference-based phenomena, particularly metonymy:

[1] literalness is distinct from the others which are all non-literal;

[2] metonymy can form into chains;

[3] metonymy always seems to occur with one of the other three; and

[4] metaphor and anomaly are the hardest to tell apart (and thus require the most extended processing to distinguish).

If the met* method is viewed in parallel processing terms then literal, metonymic, metaphorical, and anomalous interpretations are all sought at the same time.

Below are four types of metonymic concepts encoded in metonymic inference rules in CS, ordered from most common (synechdoche) to least, together with example sentences in which they occur. The first two concepts, PART FOR WHOLE and PROPERTY FOR WHOLE are source-driven; the others are target-driven.

PART FOR WHOLE

(8) "Arthur Ashe is *black*" (= skin coloured black).

(9) "Arthur Ashe is *white*" (= skin coloured white).

In (8) and (9), Arthur Ashe's skin, part of his body, is coloured black (white).

PROPERTY FOR WHOLE

(10) "John McEnroe is *yellow*" (= limited in bravery).

(11) "Natalia Zvereva is *green*" (= limited in experience).

In (10), for example, John McEnroe is limited with respect to his bravery, a property possessed by humans and other animals.

CONTAINER FOR CONTENTS

(3) "Dave drank the *glasses*" (= the potable liquid in the glasses).

(12) "Denise drank the *bottle*" (= the potable liquid from the bottle).

ARTIST FOR ARTFORM

(13) "Anne reads *Steinbeck*" (= the writings of Steinbeck).

(14) "Ted played *Bach*" (= the music of Bach).

In (14), for example, Ted does not literally play the composer Bach – he plays music composed by him.

CO-AGENT FOR ACTIVITY

(15) "Ashe played *McEnroe*" (= tennis with McEnroe).

If the sense of 'play' here is understood as "to play a professional sport" then Ashe does not literally play McEnroe because he is not a professional sport. Instead, Ashe plays a sport with McEnroe – most probably the sport of tennis.

A further step distinguishes metaphorical from anomalous semantic relations. Recall that a metaphorical relation contains a relevant analogy, as in (1) and (6), while an anomalous relation does not, as in (7). A relevant analogy is found by matching the relevant cell from the source sense-frame with one of the cells from the target sense-frame. If the match of cells is composed of a set of sister network paths between corresponding word senses in those cells then this is interpreted as analogical and hence indicative of a metaphorical relation. Any other match of cells is interpreted as not analogical and thus an anomalous semantic relation is recognized (see Fass, 1986, 1987).

The third component of CS is the semantic vector which is a form of representation, like the sense-frame; but sense-frames represent knowledge, whereas semantic vectors represent coherence. Semantic vectors are therefore a kind of *coherence representation*. A semantic vector is a data structure that contains nested labels and ordered arrays structured by a simple dependency syntax. The labels form into sets. The outer sets of labels indicate the application of the three kinds of lexical semantic constraints. The outermost set of labels is 'preference' and 'assertion'. The middle set is 'relevant' and 'non-relevant'. The innermost set is the kind of mapping used: 'network path' and 'cell matches'. The nesting of labels shows the order in which each source of knowledge was introduced. The ordered arrays represent the subkinds of each kind of mapping. Five-column arrays are for the five network paths; seven-column arrays are for the seven types of cell match. Each column contains a positive number which shows the number of occurrences of a particular network path or cell match.

The fourth component of CS is the process of screening. During analysis of a sentence constituent, a semantic vector is created for every pairwise combination of word senses. These word sense combinations are called *semantic readings* or simply "readings." Each reading has an associated semantic vector. Screening chooses between two semantic vectors and hence their attached semantic readings. Rank orderings among semantic relations are applied. In the event of a tie, a measure of conceptual similarity is used.

The ranking of semantic relations aims to achieve the most coherent possible interpretation of a reading. The class of preference-based semantic relations takes precedence over the class of assertion-based semantic relations for lexical disambiguation. The rank order among preference-based semantic relations

is

literal –> metaphorical –> anomalous.

If the semantic vectors are still tied then the measure of conceptual similarity is employed. This measure was initially developed to test a claim by Tourangeau and Sternberg (1982) about the aptness of a metaphor. They contend that aptness is a function of the distance between the conceptual domains of the source and target involved: the claim is that the more distant the domains, the better the metaphor. This is discussed further in section 3.3. The conceptual similarity measure is also used for lexical ambiguity resolution (see Fass, 1988c).

2.3. The Meta5 Program

CS has been implemented in the meta5 natural language program. The meta5 program is written in Quintus Prolog and consists of a lexicon holding the sense-frames of just over 500 word senses, a small grammar, and semantic routines that embody collation and screening, the two processes of CS. The program is syntax-driven, a form of control carried over from the structure of earlier programs by Boguraev (1979) and Huang (1985) on which meta5 is based. Meta5 analyzes sentences, discriminates the seven kinds of semantic relation between pairs of word senses in those sentences (i.e., the program recognizes metonymies, metaphors, and so on), and resolves any lexical ambiguity in those sentences. Meta5 analyzes all the sentences given in this paper, except (4) and (5) because they contain syntactic constructions that cannot be handled by meta5's present grammar.

Below are simplified versions of some of the metonymic inference rules used in meta5. A metonymic concept contains three key elements: the conceptual relationship, the direction of inference, and replacement of the source or target. The inference rules contain all three key elements. The rules, though written in a prolog-like format, assume no knowledge of Prolog on the part of the reader and fit with the role of metonymy shown in Figure 3, the flow-chart of preference-based semantic relations.

Each metonymic inference rule has a left hand side and a right hand side. The left hand side is the top most statement and is of the form *metonymic_inference_rule*(Source, Target). The right hand side consists of the remaining statements. Except for the bottommost one, these statements represent the conceptual relationship and the direction of inference. The bottommost statement controls the substitution of the discovered metonym for either the source or target; the statement is a call to find a new sense-network path.

• PROPERTY FOR WHOLE: source-driven.

metonymic_inference_rule(Source, Target):-

find_cell(Source, [Whole, have1, it1]), [1]
find_sense_network_path(Whole, Target). [2]

This rule represents PROPERTY FOR WHOLE, which is source-driven. The source is a property possessed by the whole in a property-whole relation, represented in statement [1]. The "whole" is the substitute metonym that replaces the source; the next sense-network path is sought between the whole and the target ([2]).

• CONTAINER FOR CONTENTS: target-driven.

metonymic_inference_rule(Source, Target):-

find_cell(Target, [it1, contain1, Contents]), [1]
find_sense_network_path(Source, Contents). [2]

This metonymic concept is target-driven. The target is the "container" in a container-contents relation ([1]). The "contents" is the substitute metonym that replaces the target. The next sense-network path is sought between the source and the contents ([2]).

• ARTIST FOR ARTFORM: target-driven.

metonymic_inference_rule(Source, Target):-

find_cell(Target, [it1, Make, Artform]),	[1]
confirm_type(create1, Make),	[2]
confirm_type(art_form1, Artform);	[3]
find_sense_network_path(Source, Artform).	[4]

Again, the inference in ARTIST FOR ARTFORM is from the target. The target is the "artist" in an artist-artform relation: the artist makes artforms ([1]). Checks are done to confirm that the making is a type of creating ([2]) and that the "artform" is a type of artform1 ([3]). The "artform" is the substitute metonym that replaces the target. A new sense-network path is computed between the source and the artform ([4]). I will now describe how meta5 recognizes some metonymies and metaphors.

(14) "Ted played *Bach*" (= the music of Bach).

In (14), there is a chain of metonymies plus a literal relation. The chain consists of ARTIST FOR ARTFORM and CONTAINER FOR CONTENTS metonymies. Both metonymic concepts are target-driven. In ARTIST FOR ARTFORM the inference is from the ARTIST (the target) to the ARTFORM (the source), so the substitute metonym replaces the target (the ARTIST) if the inference is successful.

```
                sf(play12,
                   [[arcs,
                     [[supertype, perform1]]],
                    [node2,
                     [[agent,
                       [preference, human_being1]],
                      [object,
                       [preference,music1]]]]]).

sf(music1,                                  sf(johann_sebastian_bach,
   [[arcs,                                     [[arcs,
     [[supertype, [sound1, art_form1]]]],        [[supertype, composer1]]],
    [node0,                                     [node0,
     [[musician1, play12, it1]]]]).              [[animacy1, dead1],
                                                  [sex1, male1],
                                                  [born1, 1685],
                                                  [died1, 1750]]]]).
```

Figure 4. Sense-frames of music1, play12, and johann_sebastian_bach.

The sense-frames of the verb-sense play12, music1 and johann_sebastian_bach are shown in Figure 4. The semantic relation results from matching the object preference of play12, which is for music, against the surface object, which is 'Bach', short for 'Johann Sebastian Bach'. The preference is the source and the surface object is the target.

We will follow what happens using the flow chart of Figure 3. (Enter the flow chart.) (Enter the uppermost statement box.) The sense-network path between the source (music1) and the target (johann_sebastian_bach) is sought.

(Enter the uppermost decision box of the the flow chart.) The path is not inclusive so metonymic inference rules are applied.

(Enter the middle decision box of the the flow chart.) The rule for ARTIST FOR ARTFORM succeeds. The discovered metonymic inference is that johann_sebastian_bach (the ARTIST) composes musical pieces (the ARTFORM). The metonymic inference is driven from the target (the ARTIST), which is johann_sebastian_bach. The successful metonymic inference is as follows: <1> johann_sebastian_bach (the ARTIST) is a composer1, <2> composers compose1 musical pieces (the ARTFORM). Two additional tests confirm <1> and <2>, which are that <3> composing is a type of creating, and <4> a musical_piece1

is a type of art_form1.

(Enter the leftmost statement box.) The original target (johann_sebastian_bach) is replaced by the substitute metonym (musical_piece1).

(Enter the uppermost statement box for a second time.) The sense-network path between the source (music1) and the new target (musical_piece1) is sought.

(Enter the uppermost decision box for a second time.) The path is not inclusive so metonymic inference rules are applied.

(Enter the middle decision box for a second time.) The rule for CONTAINER FOR CONTENTS succeeds. The successful inference is that a musical_piece1 (the CONTAINER) contains music1 (the CONTENTS).

(Enter the leftmost statement box for a second time.) The direction of inference in the CONTAINER FOR CONTENTS metonymic concept is from the target (the CONTAINER) towards the source (the CONTENTS), so the target (the CONTAINER) is replaced by the substitute metonym when an inference is successful. Hence in our example, the target (musical_piece1) is again replaced by a substitute metonym (music1). The implicit source, which is music1, the object preference of play12, remains unchanged.

(Enter the uppermost statement box for a third time.) The sense-network path between the source (music1) and the latest target (music1) is sought.

(Re-enter the uppermost decision box a third time.) The path is inclusive so a literal relation is found, that music1 is a type of music1.

(Exit the the flow chart.) The processing of the preference-based semantic relation(s) between play12, and its preference for music1, and johann_sebastian_bach is completed. After an initial preference violation (Johann Sebastian Bach is not a kind of music), the semantic relation found was an ARTIST FOR ARTFORM metonymic relation (that johann_sebastian_bach composes musical pieces) followed by a CONTAINER FOR CONTENTS metonymic relation (that musical pieces contain music) followed by a literal relation (that music is music).

(6) "The car drank gasoline."

There is a metaphorical relation between car1 and the verb sense drink1 in (6). The source is drink1, whose agent preference is animal1, and the target is car1 (see Figure 5).

```
                sf(drink1,
                  [[arcs,
                    [[supertype, [ingest1, expend1]]]],
                   [node2,
                    [[agent,
                      [preference,  animal1]],
                     [object,
                      [preference, drink1]]]]]).

sf(car1,                                sf(animal1,
   [[arcs,                                 [[arcs,
     [[supertype, motor_vehicle1]]],         [[supertype,  organism1]]],
    [node0,                                 [node0,
     [[it1, carry1, passenger1]]]]).         [[biology1,    animal1],
                                              [it1, drink1, drink1],
                                              [it1, eat1, food1]]]]).
```

Figure 5. Sense-frames of car1, drink1 (verb), and animal1.

A metaphorical relation is sought after failing to find either an inclusive network path or a metonymic inference between animal1 and car1, hence the network path between animal1 and car1 must be exclusive. The network path found is an estranged one.

Relevant cell of animal1	Cells of car1
[animal1, drink1, drink1]	[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [animacy1, nonliving1], [car1, roll1, [on3, land1]], [composition1, steel1], [driver1, drive1, car1], [car1, have1, [4, wheel1]], [car1, have1, engine1], **[car1, use2, gasoline1],** [car1, carry1, passenger1]]

Figure 6. Match of relevant cell from animal1 against cells of car1.

The second stage is the match between the relevant cell of animal1 and the cells of car1. As with the previous example, the same sense of the main sentence verb is relevant, drink1. The relevant cell of animal1 is again [animal1, drink1, drink1], which is matched against the cells of car1 (see Figure 6). A sister match is found between [animal1, drink1, drink1] and [car1, use2, gasoline1] from car1.

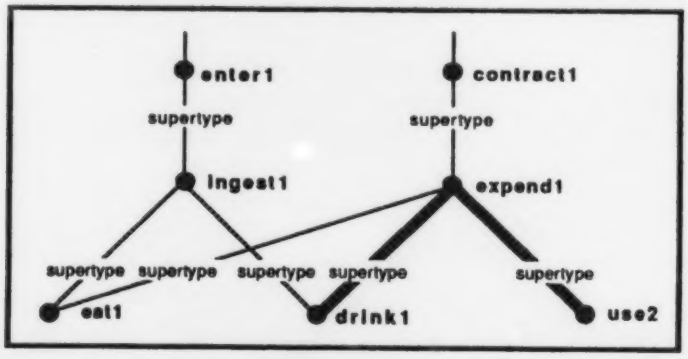

Figure 7. Sense-network path between drink1 (verb) and use2 (verb).

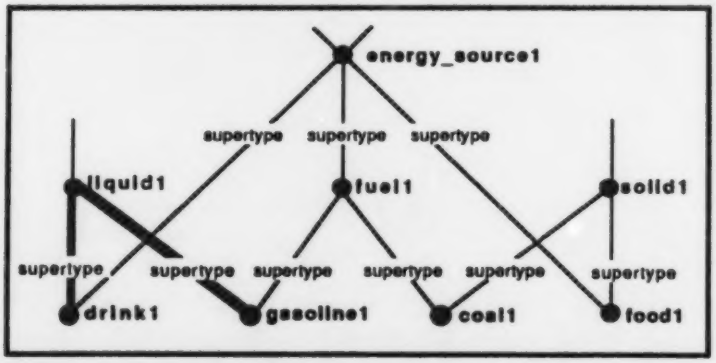

Figure 8. Sense-network path between gasoline1 and drink1 (noun).

The sister match is comprised of two sister paths found in the sense-network. The first sister path is between the verb senses drink1 and use2, which are both types of expending (Figure 7). The second path is between the noun senses drink1 and gasoline1 which are both types of liquid (Figure 8). The effect of the network paths is to establish correspondences between the two cells such that an analogy is "discovered"

that animals drink potable liquids as cars use gasoline. Note that, like Gentner's (1983) systematicity principle, the correspondences found are structural and independent of the content of the word senses they connect. Note also that the two cells have an underlying similarity or "ground" (Richards, 1936) in that both refer to the expenditure of liquids. This second stage of finding a relevant analogy seems the crucial one in metaphor recognition.

Non-relevant cells of animal1	Non-relevant cells of car1	Cell matches
[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1],	[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1],	3 same cell matches
[animacy1, living1], [composition1, flesh1],	[animacy1, nonliving1], [composition1, steel1],	2 sister cell matches
[animal1, eat1, food1], [biology1, animal1]]		2 distinctive cells of animal1
	[car1, roll1, [on3, land1]], [driver1, drive1, car1], [car1, have1, [4, wheel1]], [car1, have1, engine1], [car1, carry1, passenger1]]	5 distinctive cells of car1

Figure 9. Matches of non-relevant cells from animal1 and car1.

Figure 9 shows the match of the non-relevant cells from animal1 and car1. The cell [car1, use2, gasoline1] has been removed. There are three inclusive cell matches as animals and cars share physical object-like properties of boundedness, three dimensions, and solidity. Two cell matches are exclusive. Animals are composed of flesh, whereas cars are composed of steel. Animals are living, whereas cars are nonliving. There are two distinctive cells of animal1 and five distinctive cells of car1. Tourangeau and Sternberg's (1982) hypothesis predicts that the greater the distance between the domains of the terms involved in a metaphor, the more apt the metaphor. The proportion of similarities (inclusive cell matches) to differences (exclusive cell matches) is 3 to 2, which is a middling distance suggesting, tentatively, an unimposing metaphor.

```
[preference,
    [[network_path,
      [0, 0, 0, 0, 1]],          | First array:
     [cell_match,                | exclusive sense-network path
       [[relevant,
          [0, 0, 1, 0, 0, 0, 10]],   | Second array:
        [non_relevant,                | analogical match of relevant cell
          [0,  3,  2, 0, 0, 2,  5]]]]]   | Third array:
                                         | matches of non-relevant cells
```

Figure 10. Semantic vector for a metaphorical semantic relation.

All of these matches made by collation are recorded in the semantic vector shown in Figure 10. The crucial elements of the metaphorical relation in (6) are the preference violation and the relevant analogy. In Figure 10, the preference violation has been recorded as the 1 in the first array and the relevant analogy is the 1 in the second array. The matches of non-relevant cells are recorded in the third array.

The "preference" label indicates that a preference has been matched (rather than an assertion). The 1 in the fifth column of the first array indicates the estranged network path found between animal1 and car1.

For the cell matches, the 1 in the third column of the second array is the analogical match of the relevant cell [animal1, drink1, drink1] and the cell [car1, use2, gasoline1]. The 10 is the ten distinctive cells of car1 that did not match [animal1, drink1, drink1]. This is the match of 12 cells, 1 from the source and 11 from the target (see Figure 6). The sum of array columns is:

$$((0+0+1+0+0)\times2) + ((0+10)\times1) = (1\times2) + (10\times1) = 12.$$

The 3 similarities, 2 differences, 2 distinctive cells of animal1 and 5 distinctive cells of car1 are the non-zero numbers of the final array. The 3 similarities are all same cell matches; the 2 differences are both sister cell matches. A total of 17 cells are matched, 7 from the source and 10 from the target (see Figure 9). The total of array columns is:

$$((0+3+2+0+0)\times2) + ((2+5)\times1) = (5\times2) + (7\times1) = 17.$$

(1) "The ship ploughed the waves."

In (1), there is a metaphorical relation between a sense of the noun 'ship' and the second sense of the verb 'plough' in meta5's lexicon. Note that 'plough', like 'drink', belongs to several parts of speech. Figure 11 shows the sense-frames for the noun sense ship1, the verb sense plough2, and the noun sense plough1 which is the instrument preference of plough2.

```
sf(plough2,
    [[arcs
     [[supertype, transfer1]]],
     [node2,
      [[instrument,
        [preference, plough1]],
        [object,
        [preference, soil1]]]]]).

sf(ship1,                          sf(plough1,
  [[arcs,                            [[arcs,
   [[supertype, watercraft1]]],       [[supertype, tool1]]],
   [node0,                            [node0,
    [[it1, carry1, shipment1],         [[farmer1, plough1, it1],
     [it1, have1, engine1]]]]).         [it1, plough2, soil1]]]]).
```

Figure 11. Sense-frames of ship1, plough1 (noun), and plough2 (verb).

In (1), meta5 matches senses of 'ship' against senses of 'plough'. When meta5 pairs ship1 with plough2, it calls upon collation to match ship1 against the noun sense plough1, the instrument preference of plough2.

First, the graph search algorithm searches the sense-network for a path between plough1 (which is the preference) and ship1 and finds an estranged network path between them, i.e., a ship is not a kind of plough, so plough2's instrument preference is violated.

Next, collation inherits down lists of cells for plough1 and ship1 from their superordinates in the sense-network (for ease of reading, it1 has again been replaced by the word senses being defined). What is relevant in the present context is the action of ploughing because (1) is about a ship *ploughing* waves. Collation then runs through the list of inherited cells for the noun sense plough1 searching for a cell that refers to the action of ploughing in the sense currently under examination by meta5, plough2. Collation finds a relevant cell [plough1, plough2, soil1] and uses its frame matching algorithm to seek a match for the cell against the list of inherited cells for ship1 (see Figure 12). The algorithm finds a match with [ship1, sail2, water2] (highlighted in Figure 12), and hence collation "discovers" a relevant analogy that both ships and

ploughs move through a medium, i.e., that ploughs plough through soil as ships sail through water.

Relevant cell of plough1	Cells of ship1
[plough1, plough2, soil1]	[[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [animacy1, nonliving1], [composition1, metal1], [ship1, use2, energy_source1], [boatman1, sail1, ship1], **[ship1, sail2, water2],** [ship1, carry1, shipment1], [ship1, have1, engine1]]

Figure 12. Match of relevant cell from plough1 with cells from ship1

Finally, collation employs the frame matching algorithm a second time to match together the remaining non-relevant cells of plough1 and ship1 (see Figure 13). The cell [ship1, sail2, water2] is removed to prevent it from being used a second time.

Non-relevant cells of plough1	Non-relevant cells of ship1	Cell matches
[[composition1, matter1],	[[composition1, metal1],	1 ancestor cell match
[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [animacy1, nonliving1],	[bounds1, distinct1], [extent1, three_dimensional1], [behaviour1, solid1], [animacy1, nonliving1],	4 same cell matches
[farmer1, plough1, plough1]]	[boatman1, sail1, ship1],	1 sister cell match
	[ship1, use2, energy_source1], [ship1, carry1, shipment1], [ship1, have1, engine1]]	3 distinctive cells of ship1

Figure 13. Matches of non-relevant cells from plough1 and ship1.

```
[preference,
   [[network_path,
     [0, 0, 0, 0, 1]],                    | First array:
    [cell_match,                          | preference violation
      [[relevant,
        [0, 0, 1, 0, 0, 0, 9]],           | Second array:
       [non_relevant,                     | relevant analogy
        [1, 4, 1, 0, 0, 0, 3]]]]]]        | Third array:
                                          | matches of remaining cells
```

Figure 14. Semantic vector for another metaphorical semantic relation.

Figure 14 shows the semantic vector produced. As with Figure 10, it shows a metaphorical relation. There is a preference violation, an estranged network path indicated by the 1 in the fifth column of the first array. There is also a relevant analogy, shown by the 1 in the third column of the second array: the analogical match of the cells [plough1, plough2, soil1] and [ship1, sail2, water2]. The second array shows that 11 cells are matched, 1 from the source and 10 from the target (check against Figure 12). The sum of the array's columns is:

$$((0+0+1+0+0)\times2) + ((0+9)\times1) = (1\times2) + (9\times1) = 11.$$

In the third array, the match of non-relevant cells, there is 1 ancestor match, 4 same matches, 1 sister match, and 3 distinctive cells of ship1. Fifteen cells are matched, 6 from the source and 9 from the target (see Figure 13). The totals are:

$$((1+4+1+0+0)\times2) + ((0+3)\times1) = (6\times2) + (3\times1) = 15.$$

Semantic vectors can represent all the semantic relations except metonymic ones. The reason is that metonymic relations, unlike the others, are not discriminated by CS in terms of only five kinds of network path and seven kinds of cell matches. Instead, they consist of combinations of network paths and specialized matches of cells which have not fallen into a regular enough pattern to be represented systematically.

2.4. Extensions

Even for those semantic dependencies investigated, the interpretation of semantic relations seems to require more complexity than has been described so far in this paper. Consider the differences between the following sentences:

(6) "The car drank gasoline."

(16) "The car drank coffee."

Intuitively, sentence (6) is metaphorical while (16) is metaphorical/anomalous. In (6), the semantic relation between 'car' and 'drink' is thought to be metaphorical, and the isolated semantic relation between just 'drink' and 'gasoline' is anomalous, but the sentence as a whole is metaphorical because it is metaphorical that cars should use up gasoline.

In (16), the semantic relation between 'car' and 'drink' is metaphorical; the semantic relation between just 'drink' and 'coffee' is literal; yet the effect of (16) as a whole is metaphorical/anomalous. The object preference of 'drink' is for a drink, i.e., a potable liquid. It seems that it is metaphorical for cars to "drink" a liquid commonly used up by cars, e.g., gasoline, but anomalous if the liquid has nothing to do with cars, e.g., coffee, as in (16).

The problem of understanding the differences between sentences (6) and (16) requires some further observations about the nature of semantic relations, principally that the differences are caused by the combinations of semantic relations found in the sentences and the *relationships* between those relations. Below is a suggestion as to how deeper semantic processing might discriminate the differences between the two sentences.

Before getting to the deeper processing, we need a better semantic vector notation. The better semantic vector notation, which developed from a discussion with Afzal Ballim, is a modification of the notation shown in section 4. The key differences are *reformulation* by rewriting the five and seven column arrays in terms of the predicate-argument notation used in the rest of semantic vectors, and *extension* by adding the domain knowledge connected by every network path and cell match.

Figure 15 shows the semantic vector in Figure 10 reformulated and extended. The advantage of vectors such as the one in Figure 15 is that they record both the ways in which cells are matched (i.e., as one of seven kinds of match) and the contents (i.e., the domain knowledge) of all cells that are matched. Recording the contents of matched cells is useful because it enables a deepened analysis of semantic relations. Such an analysis is needed to detect the differences between (6) and (16).

In the description of CS in section 3, collation discriminates the one or more semantic relations in each semantic dependency, but treats the semantic relations in one dependency as isolated from and unaffected by the semantic relations in another dependency. What is needed is extra processing that interprets the semantic relation(s) in a later dependency with respect to the semantic relation(s) established in an earlier one. This processing matches the domain knowledge in semantic vectors, i.e., this processing is a

comparison of coherence representations.

```
[preference,
 [[network_path,
     [estranged,
        [1, [animal1, car1]]]],
  [cell_match,
   [[relevant,
     [[sister,
        [1, [[[animal1, drink1, drink1], [car1, use2, gasoline1]]]],
      [distinctive_target,
        [10, [[bounds1, distinct1], [extent1, three_dimensional1],
              [behaviour1, solid1], [composition1, metal1],
              [animacy1, nonliving1], [car1, roll1, [on3, land1]],
              [driver1, drive1, car1], [car1, have1, [4, wheel1]],
              [car1, have1, engine1], [car1, carry1, passenger1]]]]],
    [non_relevant,
      [[same,
        [3, [[[bounds1, distinct1], [bounds1, distinct1]],
             [[extent1, three_dimensional1], [extent1, three_dimensional1]],
             [[behaviour1, solid1], [behaviour1, solid1]]]],
       [sister,
        [2, [[[composition1, flesh1], [composition1, metal1]],
             [[animacy1, living1], [animacy1, nonliving1]]]],
       [distinctive_source,
        [2, [[animal1, eat1, food1], [biology1, animal1]]]],
       [distinctive_target,
        [5, [[car1, roll1, [on3, land1]], [driver1, drive1, car1],
             [car1, have1, [4, wheel1]], [car1, have1, engine1],
             [car1, carry1, passenger1]]]]]]]]]]
```

Figure 15. Reformulated and extended version of Figure 10.

In sentences such as (6) and (16) there are two key semantic dependencies. The first one is between the subject noun and the verb; the second is between the verb and object noun. In each dependency, the source is the verb via its agent and object preferences, and the targets are the nouns. Semantic relations are found for each dependency. One way to detect the difference between metaphorical sentences such as (6) and metaphorical/anomalous ones such as (16) is in each sentence to consult the semantic vectors produced in its two main semantic dependencies and compare the matches of the relevant cells that are found by collation.

```
[preference,
   [cell_match,
      [relevant,
         [sister,
            [1, [[animal1, drink1, drink1], [car1, use2, gasoline1]]]]]]]
```

Figure 16. Vector statement of match of relevant cell from animal1 against car1.

Let us go through such an analysis using CS, starting with the first semantic dependency between subject noun and verb. In this semantic dependency in both (6) and (16), a relevant analogy is discovered as part of a metaphorical relation between the target car1 and animal1, the agent preference of the source drink1. The semantic vector in Figure 15 records the two cells that figure in that relevant analogy. Figure 16 shows the same information from the semantic vector but written as a statement.

When the second semantic dependency is analyzed in (6), the target is gasoline1 and is matched against the noun sense drink1, the object preference of the source drink1 (the verb sense). A semantic vector is produced. The relevant cell found in the noun sense drink1 is [animal1, drink1, drink1]. Its match against [vehicle1, use2, gasoline1], a cell from gasoline1, is shown in the vector statement in Figure 17. The match is a sister match, indicating a relevant analogy.

Now this is peculiar because "drinking gasoline" is anomalous yet a relevant analogy has been found and this paper has argued that relevant analogies are special to metaphorical relations. One possible explanation is that differences exist between the recognition of metaphorical relations that concern agents and metaphorical relations that concern objects and other case roles. It may be that metaphorical relations are indicated by a relevant analogy, but only *in selected circumstances*. This needs further investigation.

```
[preference,
   [cell_match,
      [relevant,
         [sister,
            [1, [[animal1, drink1, drink1], [vehicle1, use2, gasoline1]]]]]]]
```

Figure 17. Vector statement of match of relevant cell from drink1 (noun) against gasoline1.

To return to the analysis of (6), what appears to be important in determining that (6) is a metaphorical sentence is the comparison of the two pairs of matched relevant cells:

[[animal1, drink1, drink1], [car1, use2, gasoline1]]

[[animal1, drink1, drink1], [vehicle1, use2, gasoline1]]

The two source cells are the same and the two target cells, [car1, use2, gasoline1] and [vehicle1, use2, gasoline1], are almost identical, indicating that the same basic analogy runs through the whole of (6), hence the sentence as a whole is metaphorical.

Now let us analyze the second semantic dependency in (16). The target is coffee1 and is again matched against drink1, the object preference of the verb sense drink1, the source. The relevant cell from the noun sense drink1 is again [animal1, drink1, drink1] which matches against [human_being1, drink1, coffee1] from the target coffee1. This time, the match is an ancestor match and hence not a relevant analogy. Figure 18 shows this match of the relevant cell as a vector statement.

```
[preference,
   [cell_match,
      [relevant,
         [ancestor,
            [1, [[animal1, drink1, drink1], [human_being1, drink1,  coffee1]]]]]]]
```

Figure 18. Vector statement of match of relevant cell from drink1 (noun) against coffee1.

Let us compare the two pairs of matched relevant cells for (16):

[[animal1, drink1, drink1], [car1, use2, gasoline1]]

[[animal1, drink1, drink1], [human_being1, drink1, coffee1]]

The two source cells are the same but the two target cells, [car1, use2, gasoline1] and [human_being1, drink1, coffee1], are very different. The reason that the sentence as a whole is metaphorical/anomalous is because of the clash between these target cells. The basic analogy of a car ingesting a liquid does not carry over from the first semantic dependency into the second. The anomalous flavour of (16) could not be detected by looking at the semantic relations in the dependencies in isolation because one semantic relation is metaphorical and the other is literal; neither relation is anomalous – the

anomaly comes from the interaction between the two relations.

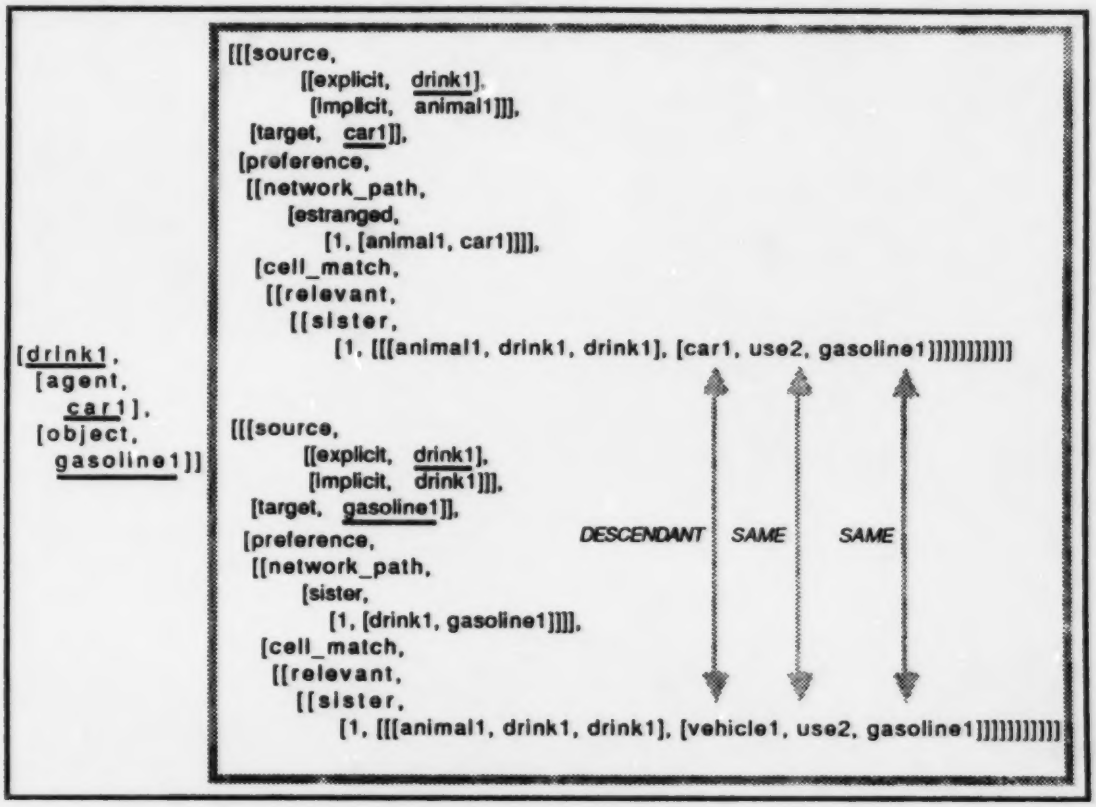

Figure 19. Sentence representation for "The car drank gasoline."

Figure 19 is a proposed representation for sentence (6). The left side of Figure 19 shows the knowledge representation part of the sentence representation: a simple case-frame based representation of (6). The right side of Figure 19, within the grey partition, is the coherence representation component of the sentence representation: abridged semantic vectors for the two main semantic dependencies in (6). The upper semantic vector is the match of the target car1 against the implicit source animal1. The lower semantic vector is the match of the target gasoline1 against the implicit source drink1, the noun sense. The upper abridged semantic vector indicates a metaphorical relation. The lower semantic vector also indicates a metaphorical relation though, as was noted earlier, "drinking gasoline" when interpreted in isolation is surely anomalous.

The underlines in Figure 19 denote pointers linking the semantic vectors to the case frame. The grey vertical arrows show that the two semantic vectors are also linked together via the matches of their relevant cells. In those matches, the arrows are sense-network paths found between the elements of the two target cells. The network paths indicated in grey that connect the two abridged semantic vectors are performing *processing of coherence representation*. The particular network paths found (indicated in italics), a descendant path and two same "paths," show that the same relevant analogy is used in both semantic relations, that both semantic relations involve a match between animals drinking potable liquids and vehicles (including cars) using gasoline, hence sentence (6) as a whole is metaphorical. Figure 19 is therefore unlike any of the coherence representations shown previously because it shows a representation of a metaphorical *sentence*, not just two isolated metaphorical relations.

Compare Figure 19 with Figure 20, a sentence representation for (16). The upper semantic vector again indicates a metaphorical relation between car1 and drink1. The lower semantic vector indicates a literal relation between drink1 and coffee1. What is important here is the match of relevant information

discovered in the two semantic relations, as indicated by the three network paths. The paths found are two estranged paths and a sister path, indicating that the relevant information found during the two semantic relations is different: in one semantic relation, information about animals drinking potable liquids is matched against *cars using gasoline*; in the other, the same information is matched against *humans beings drinking coffee*; but cars using gasoline and human beings drinking coffee are quite different, hence sentence (16) is anomalous overall.

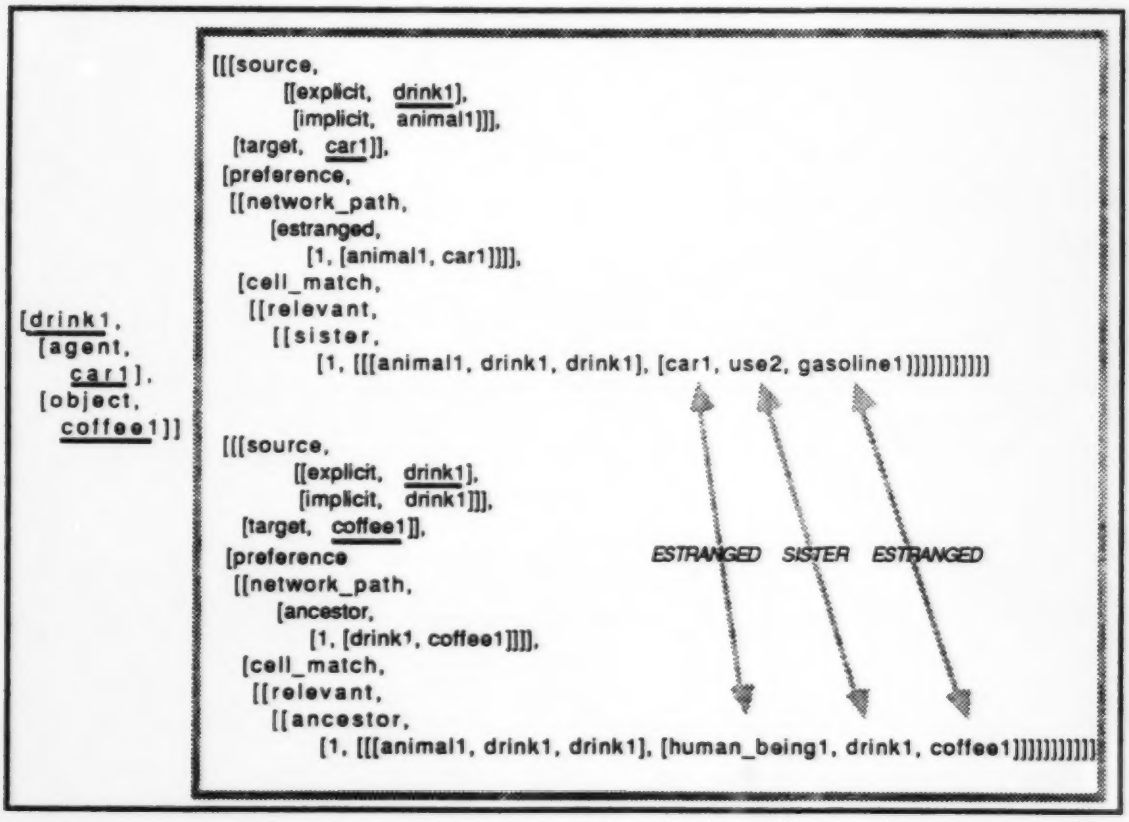

Figure 20. Sentence representation for "The car drank coffee."

Note that in Figures 19 and 20, the coherence representation part of the sentence representation is much larger than the knowledge representation part. The detailed "world knowledge" about car1, the verb sense drink1, gasoline1, and coffee1 are all on right side. It is interesting to contrast the figures with CD diagrams, e.g., those in Schank (1973), because what the figures present is a quantity of world knowledge, but only the world knowledge needed to discriminate the semantic relations in (6) and (16) rather than the much larger and seemingly unlimited amounts that appear in CD diagrams.

3. Non-Computational Views of Metonymy and Metaphor

Despite the large literature on metaphor and metonymy, the two phenomena have been on the whole poorly defined and remain vague notions. This is especially true of the non-computational literature which includes philosophical, linguistic and psychological approaches which, with the exception of, e.g., Levin's work in linguistics, lack detailed theoretical treatments of either phenomena. The two phenomena are so poorly understood that widely divergent views exist about the relationship between metaphor and metonymy. Those views are described in section 3.1. To my knowledge, there are no theories about metonymy of any depth, but some general theories of metaphor exist, which are described in section 3.2. In section 3.3., there is some discussion of the relationship of CS and the met* method to those parts of the non-computational literature reviewed.

3.1. The Relationship Between Metonymy and Metaphor

Both metonymy and metaphor have been identified as central to the development of new word senses, and hence to language change (see, e.g., Stern, 1931; Waldron, 1967). Some of the best examples of the differences between the two phenomena come from data used in studies of metonymic and metaphorical effects on language change. Nevertheless, there are widely differing views on which phenomenon is the more important. Some argue that metaphor is a kind of metonymy, others propose that metonymy is a kind of metaphor, while still others suggest that they are very different (see Fass, 1988c, p. 177). Among the third group, a common description of the difference is that metonymy is founded on contiguity whereas metaphor is based on similarity (cf. Jakobsen & Halle, 1956, pp. 76-77; Ullmann, 1962, pp. 212-220). Lakoff and Johnson (1980) offer a related description of the difference, that metaphor is

> "principally a way of conceiving of one thing in terms of another, and its primary function is understanding" (Ibid., pp. 36-37)

whereas metonymy

> "has primarily a referential function, that is, it allows us to use one entity to *stand for* another" (Ibid., their italics)

though it has a role in understanding because it focusses on certain aspects of what is being referred to.

3.2. Theories of Metaphor

The four views discussed are the selection restriction violation view, the comparison view, the interactive view, and the conventional metaphor view. Following Tourangeau and Sternberg, these views will be discussed with respect to two issues: *recognition* of metaphors and determining the *aptness* of metaphors. The difference here is that it is one issue to distinguish a metaphor from other phenomena, but it is another issue to say how good a metaphor is.

The selection restriction violation view has also been called "the semantic deviance view" (Johnson, 1980, p. 50) and "the anomaly view" (Tourangeau & Sternberg, 1982, p. 211). Johnson (1980, p. 50) describes this as a common view among linguists; Tourangeau and Sternberg (1982, p. 209) list the following people as holders of this view: Beardsley (1962), Bickerton (1969), Campbell (1975), Guenther (1975), Percy (1954), Van Dijk (1975), Wheelwright (1962). To this one might add Levin (1977). Johnson (1980) describes this view as where:

> "metaphor constitutes a violation of selection restriction rules within a given context, where the fact of this violation is supposed to explain the semantic tension one experiences in comprehending any live metaphor" (p. 50).

The extreme version of this view, as Ortony (1980, p. 70) points out, is Matthews' (1971) statement that:

> "the presence of a selection restriction violation is "a necessary and sufficient condition for the distinguishing of metaphor from non-metaphor"."

In this extreme view, according to Tourangeau and Sternberg (1982, pp. 209-210), *dissimilarity* rather than similarity is taken to be central in understanding the nature of metaphors, and theorists differ on the exact effect that results from the violation and on how violated selection restrictions should be handled. Proposals include abstraction, for example, Chomsky (1964) handles a violation by

> "finding a higher-level category to which the problematic item belongs. At this more abstract level, the violation disappears" (Tourangeau & Sternberg, 1982, p. 210).

Another proposal is special rules for handling metaphorical language (e.g., Levin, 1977).

On aptness, Tourangeau and Sternberg say:

> "we might predict that, since clarity suffers as a function of distance between tenor and vehicle, the best metaphors involve the least distance" (p. 210).

Ortony (1980) distinguishes two problems with the selection restrictions view. The one problem is that if a selection restriction is a necessary and sufficient condition for recognizing a metaphor then

> "there would seem to be no way to distinguish metaphor from semantic anomaly or falsehood" (p. 71).

Matthews' answer is to assert that a speaker must intend to use language metaphorically, i.e., that it is a matter of speaker intention.

(17) "Sierra Leone is the largest city in Italy."

But, says Ortony, one who utters (17) "is not likely to be speaking metaphorically" (p. 71). The problem, more generally, is that

> "Matthews argues that the selection restriction violation must be from the speaker's point of view, but this places metaphor essentially in the mind of the speaker, and without the introduction of some kind of pragmatic analysis, there it has to stay" (p. 71).

The second problem is that there are perfectly well-formed sentences that are metaphorical and yet contain no selection restriction violations.

> "They are sentences that demand a metaphorical interpretation in some contexts, and a literal interpretation in others" (p. 70).

(18) "The rock is becoming brittle with age" (Reddy, 1969, p. 242).

In (18), there is a literal interpretation and a metaphorical one when uttered about a decrepit professor emeritus. Johnson (1980, p. 50) notes the same problem, describing it as a problem of elevating a condition that occasionally holds (i.e., a selection restriction violation) into a necessary condition of metaphor.

(19) "I was a morsel for a monarch" (Shakespeare, *Antony and Cleopatra*, I/v/30).

Johnson agrees with Loewenberg's (1975, p. 315) observation that the interpretation of (19) might be literally true if the King was a cannibal. (Unfortunately, (19) cannnot be figuratively true because if "I" has been a morsel for a monarch then I cannot be alive to report this in the past tense. I would agree with Loewenberg and Johnson if the first word of (19) was 'She' or 'He'.) Johnson also notes what Cohen calls "twice-true metaphors" such as (20).

(20) "Idi Amin is an animal" (Johnson, 1980, p. 51).

Another example of a twice-true metaphor is (21).

(21) "No man is an Island" (John Donne, *Meditations XVII*).

The second problem, then, is that viewing metaphor in terms of only selection restriction violations ignores the influence of context. As Johnson puts it:

> "It does seem clear that our apprehension of an utterance as metaphorical relies upon our awareness of the total speech situation in which it occurs. We seem to interpret an utterance metaphorically when to do so makes sense of more aspects of the total context than if the sentence is read literally. Consider the simple case of the sentence *All men are animals* as uttered by Professor X to an introductory biology class and as uttered later by one of his female students to her roommate upon returning from a date. In the latter instance the roommate understands the utterance as metaphorical, because to do so makes better sense of the utterance *taken in its total context* .... Thus it appears that an adequate account of the identification and comprehension of metaphor must consider the role of the extra-sentential and even the extra-linguistic context of an utterance" (p. 51).

In a similar way, Ortony (1980) suggests that metaphor should be thought of as

> "*contextually* anomalous. This means that a literal interpretation of the expression, be it a word, phrase, sentence, or an even larger unit of text, fails to fit the context" (p. 73, his italics),

so whether or not a sentence is a metaphor depends upon the context in which it is used:

> "*if* something is a metaphor then it will be contextually anomalous if interpreted literally .... Insofar as the violation of selection restrictions can be interpreted in terms of semantic incompatibilities at the lexical level, such violations may sometimes be the basis of the contextual anomaly. But there can be other reasons too, so that selection restriction violations, or, to use the theoretically more neutral description, lexical level semantic incompatibilities, are not the

only causes of contextual anomaly" (p. 74).

According to Tourangeau and Sternberg (1982, p. 211), comparison theorists take *similarity* to be central in understanding the nature of metaphors:

> "According to the comparison view, a metaphor is a comparison in which one term (the *tenor* or subject of the comparison) is asserted to bear a partial resemblance (the *ground* of the comparison) to something else (the *vehicle*), the resemblance being insufficient to sustain a literal comparison. As with any comparison, there is always some residual dissimilarity (the *tension*) between the terms involved in the comparison, but comparison theorists tend not to emphasize this dissimilarity" (p. 205).

(22) "Prudence is a rich, ugly old maid, courted by Incapacity" (William Blake, *Proverbs of Hell*).

Hence, in (22), an aphorism of Blake's,

> "the first part ... could ... be paraphrased: 'Prudence is like a rich, ugly, old maid, in that it is vain, unattractive, cautious, conservative, infertile, slow to act, etc.'. Prudence is the tenor or subject of the comparison: vanity, sterility, age, conservatism, and so on, form the ground of the metaphor – the basis for the comparison" (p. 205).

Aristotle proposed the first comparison theory and

> "suggested several principles for finding the ground of the metaphor. His principles reduce to finding a category to which the tenor and vehicle belong or to constructing an analogy involving them" (Ibid.).

According to Tourangeau and Sternberg, the first principle has proven "very popular" and has been

> "rivalled only by the closely related view that the tenor and vehicle share some set of features and that these shared features form the ground or basis for the metaphor. Almost all theorists – including many who reject other aspects of the comparison view – assume that the ground of a metaphor consists of common category memberships, or a set of features, shared by tenor and vehicle (Basso, 1976; Bickerton, 1969; Guenther, 1975; Malgady and Johnson, 1976; Ortony, 1979b; Tversky, 1977; Van Dijk, 1975)" (p. 206).

On aptness: the view is that "the better the comparison, the better the metaphor," (p. 207) though "with too much similarity, the metaphor seems to suffer" (p. 208). So,

> "Comparison theorists, as a result, sometimes predict either that the metaphors improve until some cutoff of similarity is reached, or that intermediate levels of similarity produce the best metaphor" (Ibid.).

Tourangeau and Sternberg list some problems with the comparison view, including the following:

> "(a) that everything has some feature or category that it shares with everything else, but we cannot combine just any two things in metaphor; (b) that the most obvious shared features are often irrelevant to a reading of the metaphor; (c) that even when the feature is relevant, it is often shared only metaphorically; ... and (e) that metaphors are novel and surprising is hard to reconcile with the idea that they rely completely on extant similarities" (pp. 226-227).

Johnson (1980) also notes problem (a) with comparison theories, pointing out that as a result they cannot account for the semantic tension between the two terms of a metaphor:

> "the comparison theory (in either form) tries to circumvent the experienced semantic strain by interpreting metaphor as nothing but a way of comparing two things to see in what respects they are alike. And since any two things are similar in some respects, this kind of theory can never explain what is interesting and important about metaphor" (p. 52).

> "Interaction theorists (Black, 1962; Hesse, 1966; Miles, 1967; Richards, 1936; Wheelwright, 1962) try to stress similarity and dissimilarity equally. Interaction theorists argue that the vehicle of a metaphor is a template for seeing the tenor in a new way. This reorganization of the tenor is necessary, because the characteristics or features of the vehicle cannot be applied directly to the tenor; the features they 'share' are often only shared metaphorically. As Black (1962) observes, the ground of a metaphor may itself be nonliteral. 'Men are wolves', in Black's example, in part because both are predators; but they are predators in sharply different

senses that may only strike us as similar when we interpret the metaphor. In Black's reading of this metaphor, we see competition in social relations as corresponding to predacity in beasts" (Tourangeau & Sternberg, 1982, pp. 212-213).

A problem with the view is that theorists have not provided much detail about the processes involved. Richards (1936), the originator of the view,

"lapses into the comparison framework when he speaks of the 'common characteristics' of tenor and vehicle as forming the ground of the metaphor" (Ibid., p. 213).

Black (1962) does make some suggestions.

"According to Black, tenor and vehicle ... each have a 'system of commonplaces' associated with them. These commonplaces are stereotypes, not necessarily definitional, not even necessarily true, just widely agreed upon. In interpreting 'man is a wolf', we 'evoke the wolf-system of related commonplaces' and are led by them 'to construct a corresponding system of implications about the principal subject (Man)' (Black, 1962, p. 41). In Black's view, then, interpretation involves not so much comparing tenor and vehicle for existing similarities, as construing them in a new way so as to create similarity between them" (Ibid.).

"Like the anomaly view, the interaction theory offers an account of the surprise and novel vision that metaphors engender. We see one thing in terms of another thing that is plainly different from it; because they are so different, to interpret the metaphor, we must see tenor and vehicle in a new way" (Ibid., p. 214).

One version of the interaction view is the domains-interaction view, set forth by Tourangeau and Sternberg. The core of their view is that

"Metaphors often involve seeing in a new way not only two particular things but the domains to which they belong as well. In interpreting 'Men are wolves', we not only see men and wolves in a new light but we also see the domain of social relations as being analogous to the domain of beasts; we see a characteristic, such as predacity, that applies within one domain, as being analogous to a characteristic, competitiveness, that applies within another domain. Metaphors can thus involve whole systems of concepts" (p. 214).

"In interpreting a metaphor, we see one concept in terms of another by construing features or dimensions that apply within the domain of the first concept (the subject or tenor of the metaphor) as somehow parallel to those that apply within the domain of the second concept (the vehicle); further, tenor and vehicle are asserted to have similar values on those corresponding dimensions. In metaphor, we see concepts in different domains as occupying analogous positions (i.e., having analogous features) within their domains" (p. 215).

"Most metaphors ... occur in a context; this context can determine the relevant domains" (p. 215).

"The domains-interaction view, in summary, asserts that (a) a metaphor involves seeing something in one domain in terms of something in another domain; (b) because features are often specific to a domain, they must be transformed, i.e., seen in a new way, if we are to find correspondences across domains; (c) since the features that structure the domains are reinterpreted or transformed by the metaphor, the whole domain, and not just a particular term in it, partakes in this conceptual 'interaction'; (d) either the context, or in default of this, the domains themselves, can provide the structure that makes salient the features or dimensions that figure in the interpretation of the metaphor; and (e) the domains involved place limitations on the manner by which features or dimensions applying within the domain of the vehicle can be altered so as to fit the tenor" (p. 217).

"To produce a metaphorical reading, we must identify or infer the two systems of concepts involved (i.e., the tenor and vehicle) and the domains from which they are drawn. In a sense, we are proposing that metaphors are analogies that include both tenor and vehicle and their different domains as terms" (p. 218).

"Central to our view is the assumption that features 'shared' by tenor and vehicle are often at best only analogous features, each limited in its application to one domain or another. Of

course, some features or dimensions are quite general, applying across the board to a number of domains" (p. 218).

> "We assume ... that some features or structures are 'mapped', or seen as analogous, across different domains .... Because tenor and vehicle belong to different domains, they cannot generally share features" (p. 227).

Tourangeau and Sternberg propose several methods for relating features across domains (p. 219), including the following five:

#1. Abstraction: "there may be some common feature or dimension crosscutting the two domains" (p. 221) or "we may abstract the features themselves" (p. 221); "the feature applying to the tenor and that applying to the vehicle may both map onto some more abstract, common feature" (p. 218), e.g., "in interpreting 'Ian Paisley is a George Wallace', we note that Paisley is anti-Catholic, that Wallace is anti-black, and that both anti-Catholicism and racism are forms of bigotry" (p. 219).

#2. Natural association: "the features or dimensions may be naturally correlated" (p. 221), "sometimes two dimensions are statistically related" (p. 219), e.g., 'the room was orange with warmth' may be related in this way.

#4. Sharing a common label.

#5. Finding a mediating dimension or feature in a third domain.

#6. "The concepts may have similar network structures" (p. 221).

> "[W]e have tacitly assumed that the meaning or structure of concepts uses semantic networks as the chief representational device" (p. 219).

Aptness is given special attention in the domains-interaction view. Two different kinds of similarity are distinguished that have different effects on aptness. The first kind of similarity is "within-domains" similarity which refers to

> "the extent to which tenor and vehicle occupy similar 'relative positions' within their respective domains" (p. 225).

The effect of within-domains similarity is the more similar the relative positions of tenor and vehicle, the more apt the metaphor.

The second kind of similarity is "between-domains" similarity which refers to

> "the degree to which the two domains themselves resemble each other" (p. 225).

The effect of between-domains similarity is that the further apart the domains, the more apt the metaphor.

The conventional metaphor view has been advanced by Lakoff and Johnson (1980). They argue that

> "the essence of metaphor is understanding and experiencing one kind of thing in terms of another" (p. 5),

typically in conceptualizing the nonphysical in terms of the physical. Central to this process of conceptualization, Lakoff and Johnson argue, are conventional metaphors, which are

> "metaphors that structure the ordinary conceptual system of our culture, which is reflected in our everyday language" (p. 139).

Conventional metaphors, also known as conceptual metaphors, are of three main kinds: orientational, ontological and structural. Orientational metaphors are mainly to do with kinds of spatial orientation like up-down, in-out, and deep-shallow. Example metaphors include MORE IS UP and HAPPY IS UP. They arise from human experience of spatial orientation and thus develop from the sort of bodies we have and the way they function in our physical environment. Though orientational metaphors are metaphorical concepts, they do not structure one concept in terms of another but rather organize a whole system of concepts in terms of one another (p. 14).

> "Most of our fundamental concepts are organized in terms of one or more spatial metaphors" (p. 17).

Ontological metaphors arise from our basic human experiences with physical objects (especially our own bodies) and substances. Some examples are THE VISUAL FIELD IS A CONTAINER, TIME IS A

SUBSTANCE, and THE MIND IS AN ENTITY.

> "Merely viewing a nonphysical thing as an entity or substance does not allow us to comprehend very much about it. But ontological metaphors can be further elaborated" (p. 27),

for example, THE MIND IS AN ENTITY has been elaborated as THE MIND IS A MACHINE ("I'm a little *rusty* today") and THE MIND IS A BRITTLE OBJECT ("He *cracked* under pressure").

Structural metaphors are "cases where one concept is metaphorically structured in terms of another" (p. 14), e.g., ARGUMENT IS WAR and TIME IS MONEY. Structural metaphors are elaborated orientational and ontological metaphors (cf. p. 61). In them, concepts which correspond to natural kinds of experience, e.g., PHYSICAL ORIENTATIONS, OBJECTS, SUBSTANCES, SEEING, WAR, JOURNEYS, MADNESS, and BUILDINGS, are used to define other concepts, also natural kinds of experience, e.g., LOVE, TIME, HAPPINESS, IDEAS, UNDERSTANDING, and ARGUMENTS (p. 118). Hence, for example,

> "The ARGUMENT IS WAR metaphor forms a systematic way of talking about the battling aspects of arguing .... Because the metaphorical concept is systematic, the language we use to talk about the concept is systematic" (p. 5).

The language of argument is literal. So in everyday language, we talk about actually winning or losing arguments, the person we are arguing with is an opponent, we attack his or her positions and defend our own, we gain and lose ground, and so on. We talk about arguments this way, Lakoff and Johnson claim, because we *conceive* of them this way:

> "Not only our conception of an argument but the way we carry it out is grounded in our knowledge and experience of physical combat" (p. 63).

Lakoff and Johnson thus claim that language is structured in terms of basic orientational and ontological metaphors which are elaborated as structural metaphors. Lakoff and Johnson also discuss extending and combining conventional metaphors in order to metaphorically structure language in new ways.

What Lakoff and Johnson fail to discuss is how metaphors in general, let alone individual metaphorical concepts, are recognized. Martin's (1988) work, discussed in section 4, has addressed this issue. Lakoff and Johnson do not talk in terms of the aptness of a metaphor though they do discuss in a relevant, related way concepts of metaphorical "coherence" and "consistency" (see, e.g., pp. 43-45).

3.3. Own View of Metaphor

This section relates the views of metonymy and metaphor just described to the met* method.

• Metonymy and metaphor

The met* method is consistent with the view that metaphor is based on similarity whereas metonymy is based on contiguity (cf. Jakobsen & Halle, 1956) though the difference between "similarity" and "contiguity," as understood in the met* method, needs some explanation because in a way both metonymy and metaphor are based on similarity between concepts, though of different kinds. The different kinds of similarity come from *what* conceptual information is related and *how*. Let us consider *what* is related first. In metaphor, an *aspect* of one concept is similar to an *aspect* of another concept; e.g., in (6), an aspect of the concept for animal, that animals drink potable liquids, is similar to an aspect of another concept, that cars use gasoline.

(23) "The ham sandwich is waiting for his check" (Lakoff & Johnson, 1980, p. 35).

However in metonymy, a whole *concept* is related to an *aspect* of another concept. For example, in (23) the metonymy is that the concept for ham sandwich is related to an aspect of another concept, for "the man who ate a ham sandwich."

Regarding *how* that conceptual information is related: in the case of metaphor, the met* method assigns a central role to finding an analogy, and an analogy between two terms is due to some underlying similarity between them (the ground), e.g., in the analogy that animals drinking potable liquids is like cars using gasoline, the underlying similarity is that both animals and cars ingest liquids. In an analogy, the relationship between aspects of two concepts is purely structural. In metonymies, however, the relationships are "knowledge-laden" connections, e.g., the spatial relationships in PART FOR WHOLE and

CONTAINER FOR CONTENTS.

So in summary, "similarity" in metaphor is understood to be based on structural relationships between aspects of concepts, whereas "contiguity" in metonymy is based on knowledge-specific relationships between a concept and an aspect of another concept. These observations, I would argue, support the view that metonymy has primarily a referential function, allowing something to *stand* for something else, especially that in a metonymy a concept is related to an aspect of another concept. The observations also support the view that metaphor's primary function is understanding, allowing something to be conceived of in terms of something else: the role of analogy is especially crucial to this function.

The met* method is related to all four views of metaphor just described because there are two separate stages in its recognition of metaphor:

[1] a contextual constraint violation, such as a preference violation – as in the selection restrictions view;

[2] a relevant analogy – related to the comparison and interaction views.

Also,

[3] the analogies found seem to fall into patterns not unlike those in conventional metaphors.

• Selection Restrictions View of Metaphor

I agree with the two criticisms pointed out in the last two sections. Regarding the first criticism, that other phenomenon occur in the presence of selection restriction violations, two other phenomena (metonymy and anomaly) were noted in section 2.

The second criticism was that metaphorical sentences exist that do not contain selection restriction violations. I agree with the not uncommon observation that metaphor requires interpretation in *context*. I agree with the gist of Ortony's remark that metaphor should be thought of as contextual anomaly, though I would prefer the term *contextual constraint violation* because anomaly is one of the phenomena detected by contextual violation. In CS, I have investigated the presence of metaphor in violations of preferences, a kind of lexical contextual constraint. Though clearly this is a small part of the picture, it is not a bad place to start since perhaps we can establish a picture of metaphor and preference violation before moving on to other contextual constraints.

Tourangeau and Sternberg suggest that aptness in the selection restrictions view is a function of "distance" between tenor and vehicle. Distance measures are widely used in NLP work on spreading activation, e.g., Quillian's and Hirst's, where distance is the number of arcs traversed in a semantic network. Numbers of arc traversals are not used as a measure in CS. However, a measure of conceptual similarity called the *distance number* is used for assessing the closeness of the conceptual domains of a source and target – an assessment very much like Tourangeau and Sternberg's within-domains similarity. The distance number for two sense-frames is the number of "inclusive" cell matches (ancestor or same) less the number of "exclusive ones" (sister, descendant, or estranged). For more details, see Fass (1988a, 1988d).

• The Comparisons View of Metaphor

Collation and the met* method have certain similarities with the comparison view of metaphor, especially in the cell matching process. The relevant analogies discovered in CS are indeed, to quote Tourangeau and Sternberg:

> "a comparison in which one term ... is asserted to bear a partial resemblance to something else."

The collation process gives quite a clear picture of the ground and tension in a metaphor. The ground is the most specific statement that subsumes both statements that figure in the analogy, e.g., [it1, ingest1, liquid1] is the ground for the analogy involving [animal1, drink1, drink1] and [car1, use2, gasoline1] (see Figures 7 and 8). Moreover, the details of the process match with both of Tourangeau and Sternberg's reductions of Aristotle's principles for finding the ground of the metaphor: (1) both terms in a metaphorical relation belong to a common category (in the example above, the common categories are it1, ingest1, and liquid1) *and* (2) an analogy is found between them.

I will reserve discussion of aptness until after comments are made on the interaction view.

The collation process also takes care of many of the problems Tourangeau and Sternberg note with the comparison view, which were quoted on page 22. Regarding problem (a), that "everything shares some feature or category ... with everything else", CS is in agreement: the only significant combination of features in a metaphor are those involved in a relevant analogy. Problem (b), that "the most obvious shared features are often irrelevant," is that the most obvious shared features are irrelevant to a metaphor, is borne out by experience with CS – for example, animals and cars share some basic physical object-like properties but these have a minor role in understanding cars drinking. The met* method bears out problem (c) that "even when a feature is relevant, it is often shared only metaphorically." Finally, problem (e) is the claim that novel metaphors cannot be based on "extant similarities," but the relevant analogies found in the met* method are not "extant" – they have to be actively discovered.

• The Interactions View of Metaphor

According to interaction theorists like Tourangeau and Sternberg or Johnson, the key difference between the interaction and comparison views is that novel similarities are "created" in the interaction view (accounting for the novelty and surprise in a metaphor) whereas only pre-existing similarities are found in the comparison view. However, I would agree that this difference is mistaken and that what interaction theorists call "interaction" – the matching process that creates novel similarities – is simply a sophisticated form of comparison. This is quite evident when one examines, for example, the methods Tourangeau and Sternberg propose for relating features across domains in their theory: like Richards (1936), the originator of the interaction theory, whom they criticize, they too lapse into the comparison framework when discussing how such features are matched. The methods they propose are subsumed within their own reduction of Aristotle's principles for finding the ground of a metaphor in the comparison approach. One of the reductions is that the two terms in a metaphor share a common category. All four of Tourangeau and Sternberg's methods listed earlier are instances of this principle. The sister network path in CS is another quite general instance of this principle; indeed, the four methods listed are four different ways of viewing the relationship between two sister terms. Methods #1, #2 and #5 are all ways of finding a common ancestor, which is what the sister network path also does. The common ancestor is an abstraction (cf. method #1), a common label (method #2), and a feature in a third domain (#5). Method #6, that concepts may have similar network structures, is a view of the relationship between two sister terms in terms of their common network structure rather than a common ancestor.

The second of the reductions is that an analogy is found, yet Tourangeau and Sternberg themselves say that:

> "In a sense, we are proposing that metaphors are analogies that include both tenor and vehicle and their different domains as terms" (1982, p. 218).

And, of course, finding an analogy is central to the met* method on CS.

With regards to aptness, within-domains similarity, the way Tourangeau and Sternberg describe it, sounds like a question of network relationships. In CS, five network relationships are distinguished, allowing some room in which to explore this kind of similarity. Between-domains similarity has been investigated in CS in terms of the match of all the non-relevant features of two concepts (see, e.g., Fass, 1988b, 1988d).

• The Conventional Metaphors View

The link that CS has to conventional metaphors is that certain types of ground seem to recur when meta5 processes metaphors, types which resemble Lakoff and Johnson's (1980) structural metaphors. Two types of ground have cropped up so far.

(24) "Time flies."

The first is a use-up-a-resource metaphor which occurs in (6) and in (24) when viewed as time (noun) flies (verb), both analyzed by meta5. Use-up-a-resource resembles structural metaphors like TIME IS A RESOURCE and LABOR IS A RESOURCE which, according to Lakoff and Johnson (1980, p. 66), both employ the simple ontological metaphors of TIME IS A SUBSTANCE and AN ACTIVITY IS A SUBSTANCE:

> "These two substance metaphors permit labor and time to be quantified – that is, measured, conceived of as being progressively "used up," and assigned monetary values; they allow us

to view time and labor as things that can be "used" for various ends."

The second type of ground is motion-through-a-medium, a type of ground discussed by Russell (1976). This appears in (7) and (25), again both analyzed by meta5.

(25) "The horse flew."

Incidentally, it is worth noting that structural metaphors have proven more amenable to the met* method than other kinds. When I began working on metaphor, I assumed that orientational and ontological metaphors would be easier to analyze than structural metaphors because they were less complex. However, the reverse has proved to be the case, probably because structural metaphors contain more specific concepts like "drink" and "plough" which are easier to represent in a network structure (like the sense-network of CS) so that analogies can be found between those concepts.

4. Computational Views of Metonymy and Metaphor

The theories of Wilks, Hobbs, Gentner, and Martin will each be briefly described, focussing on their treatment of metaphor and metonymy, and how they distinguish them from literalness and anomaly. Space does not permit discussion of work on metaphor by, e.g., Russell (1976), Weiner (1984, 1985), and Indurkhya (1985, 1987).

Preference Semantics addresses the recognition of literal, metaphorical and anomalous relations. The theory of metaphor in Preference Semantics consists of a selection restrictions view and a comparison view. In the theory, information about word senses is contained in knowledge structures called "semantic formulas." An algorithm matches pairs of semantic formulas seeking satisfied or violated preferences between them. A satisfied preference indicates a literal semantic relation; a violated preference indicates either a metaphorical or anomalous one. This part of the theory is implemented in a machine translation system (see Wilks, 1973).

To distinguish metaphorical from anomalous relations, a different knowledge structure and a second algorithm are used. The algorithm, called projection, operates on a knowledge structure, called a pseudo-text, that contains lists of templates (a further kind of knowledge structure) linked by case ties. A brief example of projection is given for (33).

(33) "My car drinks gasoline" (Wilks, 1978).

Projection operates only on preference violations. The best representation of (33) contains a preference violation, so projection is used. The algorithm compares the template representation for the sentence

[my+car drink gasoline]

against templates from the pseudo-text of 'car' seeking "the closest match," and selects [ICengine (USE) #liquid]. (USE) is projected onto drink in the sentence representation which becomes

[my+car use gasoline]

Hobbs' goal has been to develop a unified process of discourse interpretation based on the drawing of appropriate inferences from a large knowledge base, which Hobbs sometimes calls "selective inferencing" (e.g., Hobbs, 1980). Selective inferencing is concerned with drawing or refraining from drawing certain inferences in a controlled fashion (cf. Hobbs 1983a, p.17).

Hobbs has applied this approach to a variety of problems concerned with discourse interpretation, including coherence (1977, 1979, 1982a, 1985), lexical ambiguity (1977, 1982b: Hobbs & Martin, 1987), metaphor (1977, 1983a, 1983b), and three "local pragmatics" phenomena: reference, metonymy, and nominal compounds (Hobbs & Martin, 1987). Hobbs (1980) outlines all of these problems in one paper. Hobbs takes an interaction view of metaphor. He argues that many problems have the same or almost the same inferencing solutions. These solutions are found via four separate semantic operations that all draw inferences from text (e.g., Hobbs, 1977, pp. 111-113).

Gentner's (1983) structure-mapping theory addresses literal similarity, analogy, abstraction and anomaly which Gentner refers to as four "kinds of comparison." An algorithm compares the semantic information from two concepts represented as sets of properties. Properties are either "attributes," one-place predicates like LARGE(x), or "relations," two-place predicates such as COLLIDE(x,y). The four kinds of comparison are distinguished by the relative proportions of attributes and relations that are matched, and the forms of mappings established between them. Mappings between relations are sought before those

between attributes. Pairs of relations are compared using the "systematicity principle" that regular structural correspondences should exist between terms occupying the same positions in those relations. Mappings are purely structural and independent of the content of the relations (i.e., the predicates). This closely resembles a comparison view of metaphor.

Jacobs (1985) and especially Martin (1985, 1988) have pursued a conventional metaphor view using Kodiak, a knowledge representation language developed by Wilensky that is a variant of Brachman's KLONE. Within Kodiak, metaphorical relationships are represented using a primitive link type called a "view." A view

> "is used to assert that ... one concept may in certain circumstances be considered as another " (Martin, 1988, p. 49).

In Martin's work, "metaphor-maps," a kind of view (Ibid., p. 54), are used to represent conventional metaphors and the conceptual information they contain. Semantic vectors (of the extended kind described in section 2.4.) and metaphor-maps complement each other. In section 2.4., I argued that vectors need to record the conceptual information involved when finding mappings between a source and target. What metaphor-maps do is freeze (some of) the conceptual information involved in particular metaphorical relations. There is some convergence here between our approaches which would be interesting to explore.

None of the above has more than partial coverage of preference-based semantic relations. Preference Semantics has no means of recognizing metonymy; nor do Gentner, Hobbs or Martin. Hobbs has no treatment of anomaly.

A related criticism concerns the extent that these theories have been implemented in computer programs. In the case of Preference Semantics, the theory described in Wilks (1978) has not been implemented, though the projection algorithm was implemented (Modiano, 1986) using some parts of CS to supply detail missing from Wilks' original specification. Moreover, the Structure-Mapping Theory has been implemented in the Structure-Mapping Engine (Falkenhainer, Forbus, & Gentner, 1986) and some examples analyzed by it but not, to my knowledge, examples of metaphor or anomaly. Hobbs' papers include a series of programs – SATE, DIANA, DIANA-2, and TACITUS – but the papers are not clear about what the programs can do. It is not clear, for example, whether any of the systems actually analyze any metaphors or metonymies. In addition, Gentner's work is too removed from considerations of natural language. Hobbs' work is presently too fragmented. To my knowledge, Hobbs has yet to produce a unified description of selective inferencing that shows *in detail* how lexical ambiguity is resolved or how the differences between metaphor, metonymy, and so on can be recognised.

5. Discussion and Review

I have tried to show that the view of metonymy and metaphor in the met* method is consistent with much of the literature on those phenomena. The met* method is consistent with the view that metonymy involves contiguity whereas metaphor involves similarity, and that metonymy has

> "primarily a referential function, that is, it allows us to use one entity to *stand for* another" (Lakoff & Johnson, 1980, pp. 36-37, italics in original),

while metaphor is

> "principally a way of conceiving of one thing in terms of another, and its primary function is understanding" (Ibid.).

Metonymy and metonymy do have much in common: both might be described as forms of "conceptual ellipsis," a shorthand way of expressing ideas. In this respect, metonymy seems quite related to anaphora in that both are mechanisms that allow one entity to refer to another entity. This similarity of function can be seen in comparing (23), which is metonymic, with (34), which is anaphoric.

(23) "The ham sandwich is waiting for his check" (= the man who ordered the ham sandwich).

(34) "He is waiting for his check" (= the male person).

The met* method in its present serial form recognizes literalness, metonymy, metaphor, and anomaly in the following order and by the following characteristics:

- Literalness – an inclusive network path (i.e., satisfied preference).
- Metonymy – a successful conceptual inference.
- Metaphor – an underlying relevant analogy.
- Anomaly – none of the above.

The above analysis also illustrates, I hope, why metonymy and metaphor are easily confused: both are non-literal and are found through the discovery of some aspect (a property) shared by the implicit source, a preference, and the target, in the above case a surface noun. The difference is (a) how that aspect is selected, (b) the operations that follow, (c) the effect those operations produce, and (d) subsequent processing.

In the case of metonymy, (a) the selected aspect forms a regular semantic relationship with a property from the target; (b) there is substitution, i.e., replacement of one concept with another; (c) hence the apparent referential function of metonymy; and (d) is unclear at present.

In the case of metaphor, (a) the selected aspect is relevant; (b) forms an analogy with another aspect from the target; and (c) the effect is of surprise discovery of similarity between the two concepts; and (d) the discovered analogy is used to unearth further similarities between the two concepts (i.e, to deepen the analogy) and to guide subsequent sentence interpretation. Moreover, the view of metaphor in CS contains elements of the selection restrictions view, the comparisons view, and the interactions view of metaphor.

It should be emphasized that the met* method has only been applied to a small set of English sentences. Metonymy interpretation has been investigated only for adjective-noun and subject-verb-object constructions; metaphor interpretation, only for the latter. I believe the best avenue for progress with the met* method is the extensions to metaphor interpretation described in section 2.4. In the meantime I am looking for sentences that contain semantic relations consisting of a metonymy (or chain of metonymies) followed by a metaphor.

(35) "Prussia invaded France in 1870."

Sentence (35), which was discussed in a group working on beliefs at the CRL (see Acknowledgements), is interesting in this respect because it lends itself to *either* a metonymic *or* a metaphorical interpretation. The problem is: what is the semantic relation between 'Prussia' and 'invade'. It is non-literal because, I would argue, 'army' is the expected agent of 'invade' and 'Prussia' is a country, not an army. What, then, is the semantic relation between 'Prussia' and 'army'? One possibility is that a chain of metonymies is involved, that the army is controlled by the government which also controls Prussia. A second possibility is that Prussia is understood metaphorically as being an animate thing that extends itself into France.

6. Acknowledgements

I would like to thank the many people at the Cognitive Studies Centre, University of Essex; the Computing Research Laboratory, New Mexico State University; and the Centre for Systems Science, Simon Fraser University, with whom I have had fruitful discussions over the years, especially those in the beliefs group at the CRL (Afzal Ballim, John Barnden, Sylvia Candelaria de Ram, and Yorick Wilks); others at the CRL (including Xiuming Huang, David Farwell, and Eric Dietrich); and colleagues in the CSS who made helpful comments on earlier drafts of this paper (Chris Groeneboer, Gary Hall, and Carl Vogel). A special word of thanks for the help given by Yorick Wilks, the director of the CRL, and Nick Cercone, the director of the CSS. I also gratefully acknowledge the financial support provided by SERC Project GR/C/68828 while at Essex, by the New Mexico State Legislature while at NMSU, and by the Advanced Systems Institute and the Centre for Systems Science while at SFU.

7. References

Basso, K. (1976). 'Wise Words' of the Western Apache: Metaphor and Semantic Theory. In K. Basso & H. Selby (Eds.) *Meaning in Anthropology*. Albuquerque, NM: New Mexico Press.

Beardsley, M. (1962). The Metaphorical Twist. *Philosophy and Phenomenological Research*, 22, pp. 293-307.

Bickerton, D. (1969). Prolegomena to a Linguistic Theory of Metaphor. *Foundations of Language*, 5, pp. 36-51.

Black, M. (1962). *Models and Metaphors*. Ithaca, NY: Cornell University Press.

Black, M. (1979). More about Metaphor. In A. Ortony (Ed.) *Metaphor and Thought*. London, England: Cambridge University Press, pp. 19-43.

Boguraev, B. K. (1979). Automatic Resolution of Linguistic Ambiguities. Technical Report No.11, University of Cambridge Computer Laboratory, Cambridge, England.

Campbell, P. (1975). Metaphor and Linguistic Theory. *Quarterly Journal of Speech*, 61, pp. 1-12.

Chomsky, N. (1964). Degrees of Grammaticalness. In J. A. Fodor & J. J. Katz (Eds.). *The Structure of Language: Readings in the Philosophy of Language*. Englewood Cliffs, NJ: Prentice-Hall, pp. 384-389.

Van Dijk, T. (1975). Formal Semantics of Metaphorical Discourse. *Poetics*, 4, pp. 173-198.

Dubois, J., F. Edeline, J.-M. Klinkenberg, P. Minguet, F. Pire, & H. Trinon (1970). *Rhètorique Gènèrale*. Paris, France: Larousse. Published as *A General Rhetoric*. translated by P. B. Burrell & E. M. Slotkin, Baltimore, MD: The John Hopkins University Press, 1981.

Falkenhainer, B., K. D. Forbus, & D. Gentner (1986). The Structure-Mapping Engine. *Proceedings of the 5th National Conference on Artificial Intelligence (AAAI-86)*, Philadelphia, PA, pp. 272-277.

Fass, D. C. (1986). Collative Semantics: An Approach to Coherence. Memorandum MCCS-86-56, Computing Research Laboratory, New Mexico State University, NM.

Fass, D. C. (1987). Semantic Relations, Metonymy, and Lexical Ambiguity Resolution : A Coherence-Based Account. In *Proceedings of the 9th Annual Cognitive Science Society Conference*, University of Washington, Seattle, WA, pp. 575-586.

Fass, D. C. (1988a). Collative Semantics: A Semantics for Natural Language Processing. (PhD thesis) Memorandum MCCS-88-118, Computing Research Laboratory, New Mexico State University, NM.

Fass, D. C. (1988b). An Account of Coherence, Semantic Relations, Metonymy, and Lexical Ambiguity Resolution. In S. L. Small, G. W. Cottrell, & M. K. Tanenhaus (Eds.). *Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology, and Artificial Intelligence*. Los Altos, CA: Morgan Kaufmann, pp. 151-178.

Fass, D. C. (1988c). Metonymy and Metaphor: What's the Difference? *Proceedings of the 12th International Conference on Computational Linguistics (COLING-88)*, Budapest, Hungary, pp. 177-181.

Fass, D. C. (1988d). Collative Semantics: A Study in the Discrimination of Meaning. Technical Report CSS/LCCR TR 88-24, Centre for Systems Science, Simon Fraser University, Burnaby, BC, Canada.

Fass, D. C. (1989a). Lexical Semantic Constraints. Technical Report CSS/LCCR TR 89-11, Centre for Systems Science, Simon Fraser University, Burnaby, BC, Canada. To appear in J. Pustejovsky (Ed.) *Semantics and the Lexicon*. Dordrecht, Netherlands: Kluwer Academic Publishers.

Fass, D. C. (1989b). Four General Representations and Processes for Use in Problem Solving. In *Proceedings of 1989 International Conference on Knowledge Based Computer Systems*, Bombay, India, December 11-13.

Genette, G. (1970). La Rhetorique Restreinte. *Communications*, 16, pp. 158-171.

Gentner, D. (1983). Structure Mapping: A Theoretical Framework for Analogy. *Cognitive Science*, 7, pp. 155-170.

Geunther, F. (1975). On the Semantics of Metaphor. *Poetics*, 4, pp. 199-220.

Harris, R. (1976). Comprehension of Metaphor: A Test of a Two-Stage Processing Model. *Bulletin of the Psychonomic Society*, 8, pp. 321-324.

Hesse, M. (1966). *Models and Analogies in Science*. Notre Dame, IN: Notre Dame University Press.

Hobbs, J. R. (1977b). Coherence and Interpretation in English Texts. *Proceedings of the 5th International Joint Conference on Artificial Intelligence (IJCAI-77)*, Cambridge, MA, pp. 110-116.

Hobbs, J. R. (1979). Coherence and Coreference. *Cognitive Science*, 3, pp. 67-90.

Hobbs, J. R. (1980). Selective Inferencing. In *Proceedings of the 3rd National Conference of the Canadian Society for Computational Studies of Intelligence (CSCSI-3)*, Victoria, BC, Canada, pp. 101-

114.

Hobbs, J. R. (1982a). Towards an Understanding of Coherence in Discourse. In W. G. Lehnert & M. H. Ringle (Eds.) *Strategies for Natural Language Processing*. Hillsdale, NJ: Erlbaum Associates, pp. 223-243.

Hobbs, J. R. (1982b). Representing Ambiguity. In *Proceedings of the 1st West Coast Conference on Formal Linguistics*, Stanford, CA, pp. 15-28.

Hobbs, Jerry R. (1983a) Metaphor Interpretation as Selective Inferencing: Cognitive Processes in Understanding Metaphor (Part 1). *Empirical Studies of the Arts*, 1, pp. 17-33.

Hobbs, Jerry R. (1983b) Metaphor Interpretation as Selective Inferencing: Cognitive Processes in Understanding Metaphor (Part 2). *Empirical Studies of the Arts*, 1, pp. 125-141.

Hobbs, J. R. (1985). On The Coherence and Structure of Discourse. Report No. CSLI-85-37, Center for the Study of Language and Information (CSLI), Stanford University, CA.

Hobbs, J. R., & P. Martin (1987). Local Pragmatics. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI-87)*, Milan, Italy, pp. 520-523.

Huang, X.-M. (1985). Machine Translation in the SDCG (Semantic Definite Clause Grammars) Formalism. *Proceedings of the Conference on Theoretical and Methodological Issues in Machine Translation of Natural Languages*, Colgate University, New York, NY, pp. 135-144.

Indurkhya, B. (1985). A Computational Theory of Metaphor Comprehension and Analogical Reasoning. Boston University Technical Report #85/001, Boston University, Boston, MA.

Indurkhya, B. (1987). Approximate Semantic Transference: A Computational Theory of Metaphors and Analogies. *Cognitive Science*, 11, (4), pp. 445-480.

Jacobs, P. S. (1985). A Knowledge-Based Approach to Language Production. (PhD thesis) Report No. UCB/CSD 86/254, Computer Science Division (EECS), University of California, Berkeley, CA.

Jakobsen, R., & M. Halle (1956). *Fundamentals of Language*. The Hague, Holland: Mouton.

Johnson, Mark (1980). A Philosophical Perspective on the Problems of Metaphor. In R. P. Honeck & R. R. Hoffman (Eds.) *Cognition and Figurative Language*. Hillsdale, NJ: Erlbaum Associates, pp. 47-68.

Katz, J. J. (1964). Analyticity and Contradiction in Natural Language. In J. A. Fodor & J. J. Katz (Eds.). *The Structure of Language: Readings in the Philosophy of Language*. Englewood Cliffs, NJ: Prentice-Hall, pp. 519-543.

Lakoff, G., & M. Johnson (1980). *Metaphors We Live By*. London, England: Chicago University Press.

Lees, R. B. (1960). The Grammar of English Nominalizations. *International Journal of American Linguistics*, 26, 3.

Levin, S. R. (1977). *The Semantics of Metaphor*. Baltimore, MD: John Hopkins University Press.

Loewenberg, I. (1975). Identifying Metaphors. *Foundations of Language*, 12, pp. 315-338.

Malgady, R., & M. Johnson (1976). Modifiers in Metaphors: Effects of Constituent Phrase Similarity on the Interpretation of Figurative Sentences. *Journal of Psycholinguistic Research*, 5, pp. 43-52.

Martin, J. H. (1985). Knowledge Acquisition though Natural Language Dialogue. *Proceedings of the 2nd Annual Conference on Artificial Intelligence Applications*, Miami, FL.

Martin, J. H. (1988). A Computational Theory of Metaphor. (PhD thesis) Report No. UCB/CSD 88/465, Computer Science Division (EECS), University of California, Berkeley, CA.

Matthews, R. J. (1971). Concerning a 'linguistic theory' of metaphor. *Foundations of Language*, 7, pp. 413-425.

Miles, J. (1967). *Style and Proportion*. Boston, MA: Little, Brown and Co.

Modiano, N. (1986). Two Procedures for Sense Projection. Unpublished ms, Computing Research Laboratory, New Mexico State University, NM.

Van Noppen, J.-P., S. De Knop, & R. Jongen (1985). *Metaphor: A Bibliography of post-1970 Publications*. Amsterdam Studies in the Theory and History of Linguistic Science (Series V: Library and Information Sources in Linguistics), Volume 17, Philadelphia, PA: John Benjamins Publishing Company.

Ortony, A. (1979). Beyond Literal Similarity. *Psychological Review*, 86, pp. 161-180.

Ortony, A. (1980). Some Psycholinguistic Aspects of Metaphor. In R. P. Honeck & R. R. Hoffman (Eds.) *Cognition and Figurative Language*. Hillsdale, NJ: Erlbaum Associates, pp. 69-83.

Percy, W. (1954). *The Message in the Bottle*. New York, NY: Farrar, Strauss, and Giroux.

Quillian, M. R. (1968). Semantic Memory. In M. Minsky (Ed.). *Semantic Information Processing*. Cambridge, MA: MIT Press, pp. 216-270.

Perrine, L. (1971). Psychological Forms of Metaphor. *College English*, 33, pp. 125-138.

Reddy, M. J. (1969). A Semantic Approach to Metaphor. In *Chicago Linguistic Society Collected Papers*, Chicago, IL: Chicago University Press, pp. 240-251.

Reddy, M. J. (1979). The Conduit Metaphor – A Case of Frame Conflict in Our Language about Language. In A. Ortony (Ed.). *Metaphor and Thought*. London, England: Cambridge University Press, pp. 284-324.

Richards, I. A. (1936). *The Philosophy of Rhetoric*. Oxford, England: Oxford University Press.

Russell, S. (1976). Computer Understanding of Metaphorically Used Verbs. *American Journal of Computational Linguistics*, Microfiche 44.

Schank, R. C. (1973). Identification of Conceptualizations Underlying Natural Language. In R. C. Schank & K. M. Colby (Eds.) *Computer Models of Thought and Language*. San Francisco, CA: W.H. Freeman, pp. 187-247.

Schank, R. C. (1975). The Structure of Episodes in Memory. In D. G. Bobrow & A. Collins (Eds.). *Representation and Understanding*. New York, NY: Academic Press, pp. 237-272.

Shibles, W. (1971). *Metaphor: An Annotated Bibliography and History*. Whitewater, WI: The Language Press.

Stern, G. (1968; first published in Sweden 1931). *Meaning and Changes of Meaning*. Bloomington, IN: Indiana University Press.

Tourangeau, R., & R. J. Sternberg (1982). Understanding and Appreciating Metaphors. *Cognition*, 11, pp. 203-244.

Tversky, A. (1977). Features of Similarity. *Psychological Review*, 84, pp. 327-352.

Ullmann, S. (1962). *Semantics: An Introduction to the Science of Meaning*. Oxford, England: Basil Blackwell & Mott Ltd.

Waldron, R. A. (1967). *Sense and Sense Development*. London, England: Andre Deutsch.

Weiner, E. J. (1984). A Knowledge Representation Approach to Understanding Metaphors. *American Journal of Computational Linguistics*, 10, pp. 1-14.

Weiner, E. J. (1985). Solving the Containment Problem for Figurative Language. *International Journal of Man-Machine Studies*, 23, pp. 527-537.

Wheelwright, P. (1962). *Metaphor and Reality*. Bloomington, IN: Indiana University Press.

Wilks, Y. A. (1973). An Artificial Intelligence Approach to Machine Translation. In R. C. Schank & K. M. Colby (Eds.). *Computer Models of Thought and Language*. San Francisco, CA: W. H. Freeman, pp. 114-151.

Wilks, Y. A. (1975a). A Preferential Pattern-Seeking Semantics for Natural Language Inference. *Artificial Intelligence*, 6, pp. 53-74.

Wilks, Y. A. (1975b). An Intelligent Analyzer and Understander for English. *Communications of the ACM*, 18, pp. 264-274.

Wilks, Y. A. (1978). Making Preferences More Active. *Artificial Intelligence*, 11, pp. 197-223.

Winston, P. H. (1980). Learning by Creatifying Transfer Frames. In P. H. Winston & R. H. Brown (Eds.) *Artificial Intelligence: An MIT Perspective, Volume 1*. Cambridge, MA: MIT Press, pp. 347-376.

Yamanashi, M. (1987). Metonymic Interpretation and Associative Processes in Natural Language. In M. Nagao (Ed.) *Language and Artificial Intelligence (Proceedings of an International Symposium on Language and Artificial Intelligence held in Kyoto, Japan, 16-21 March 1986)*. Elsevier Science Publishers, B.V., (North Holland), pp. 77-86.

Centre for Systems Science

• Simon Fraser University established the **Centre for Systems Science** as a provincial centre for advanced study. The objective of CSS is twofold: coordinate and combine advanced systems research at SFU, and encourage application of research results within industry in order to further the technological development of British Columbia.

• CSS has assembled three research groups to streamline its efforts. The orientation of these groups is as follows: intelligent systems, microelectronics, and computer and communications systems. Because of shared research interests, this latter group merges with LCCR (**Laboratory for Computer and Communications Research**) for the publication of technical reports.

• CSS's three groups involve over 80 skilled faculty researchers principally from SFU's Schools of Engineering Science and Computing Science, but also draw substantially on Mathematics, Physics, Geography, Philosophy, Linguistics, Psychology, Communication, and Kinesiology.